

SinGRAF: Learning a 3D Generative Radiance Field for a Single Scene

Minjung Son^{*1,2} Jeong Joon Park^{*2} Leonidas Guibas² Gordon Wetzstein²
¹Samsung Advanced Institute of Technology (SAIT) ²Stanford University
 minjungson@samsung.com, {jjpark3d, gordon.wetzstein, guibas}@stanford.edu

Abstract

Generative models have shown great promise in synthesizing photorealistic 3D objects, but they require large amounts of training data. We introduce SinGRAF, a 3D-aware generative model that is trained with a few input images of a single scene. Once trained, SinGRAF generates different realizations of this 3D scene that preserve the appearance of the input while varying scene layout. For this purpose, we build on recent progress in 3D GAN architectures and introduce a novel progressive-scale patch discrimination approach during training. With several experiments, we demonstrate that the results produced by SinGRAF outperform the closest related works in both quality and diversity by a large margin.

1. Introduction

Creating a new 3D asset is a laborious task, which often requires manual design of triangle meshes, texture maps, and object placements. As such, numerous methods were proposed to automatically create diverse and realistic variations of existing 3D assets. For example, procedural modeling techniques [11, 27] produce variations in 3D assets given predefined rules and grammars, and example-based modeling methods [13, 21] combine different 3D components to generate new ones.

With our work, we propose a different, generative strategy that is able to create realistic variations of a single 3D scene from a small number of photographs. Unlike existing 3D generative models, which typically require 3D assets as input [13, 51], our approach only takes a set of unposed images as input and outputs a generative model of a single 3D scene, represented as a neural radiance field [31].

Our method, dubbed SinGRAF, builds on recent progress in unconditional 3D-aware GANs [5, 40] that train generative radiance fields from a set of single-view images. However, directly applying these 3D GANs to our problem

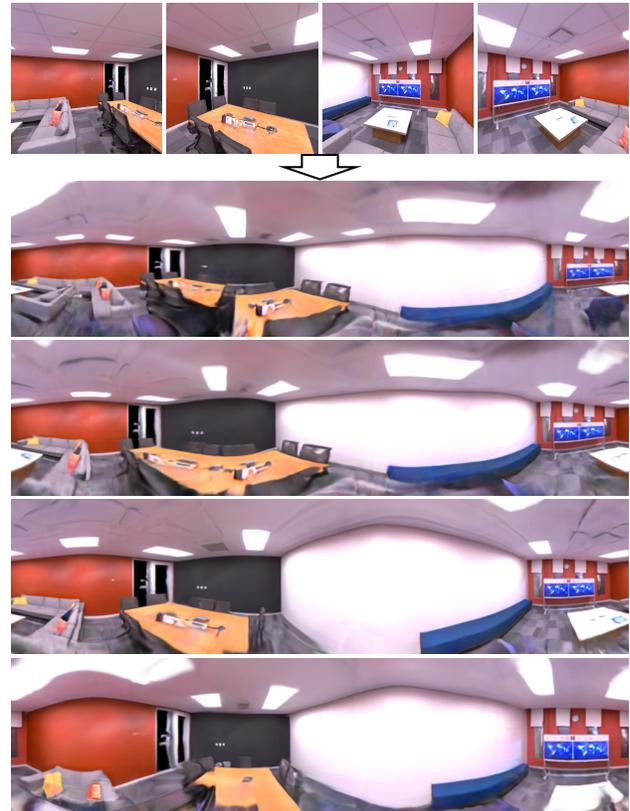


Figure 1. SinGRAF generates different plausible realizations of a single 3D scene from a few unposed input images of that scene. In this example, i.e., the “office_3” scene, we use 100 input images, four of which are shown in the top row. Next, we visualize four realizations of the 3D scene as panoramas, rendered using the generated neural radiance fields. Note the variations in scene layout, including chairs, tables, lamps, and other parts, while staying faithful to the structure and style of the input images.

is challenging, because they typically require a large training set of diverse images and often limit their optimal operating ranges to objects, rather than entire scenes. SinGRAF makes a first attempt to train a 3D generative radiance field for individual indoor 3D scenes, creating realistic 3D variations in scene layout from unposed 2D images.

^{*}Equal contribution.

Project page: computationalimaging.org/publications/singraf/

Intuitively, our method is supervised to capture the internal statistics of image patches at various scales and generate 3D scenes whose patch-based projections follow the input image statistics. At the core of our method lies continuous-scale patch-based adversarial [14] training. Our radiance fields are represented as triplane feature maps [4, 43] produced by a StyleGAN2 [22] generator. We volume-render our generated scenes from randomly sampled cameras with *varying fields of view*, to simulate the appearance of image patches at various scales. A scale-aware discriminator is then used to compute an adversarial loss to the real and generated 2D patches to enforce realistic patch distributions across all sampled views. Notably, our design of continuous-scale patch-based generator and discriminator allows patch-level adversarial training without expensive hierarchical training [41, 51, 54]. During the training, we find applying perspective augmentations to the image patches and optimizing the camera sampling distribution to be important for high-quality scene generation.

The resulting system is able to create plausible 3D variations of a given scene trained only from a set of unposed 2D images of that scene. We demonstrate our method on two challenging indoor datasets of Replica [47] and Matterport3D [6] as well as a captured outdoor scene. We evaluate SinGRAF against the state-of-the-art 3D scene generation methods, demonstrating its unique ability to induce realistic and diverse 3D generations.

2. Related Work

Synthesis from 3D Supervision. A large body of prior work aims at creating variations of scenes or objects. Procedural modeling approaches [11, 27, 33, 36] are widely used for auto-generating repetitive scenes such as terrains, buildings, or plants. These methods typically require manually designing the rules and grammars to procedurally add new 3D elements. Example-based methods [13, 21, 53] aim at extracting patterns from 3D asset examples to synthesize new models. This line of data-driven approaches learns how to mix and match different components to create a plausible 3D asset. Similarly, scene synthesis techniques [12, 38, 49] learn a distribution of plausible object arrangements from professionally designed scene datasets. All of these methods require datasets of 3D assets, part segmentation, or object arrangement designs, which are expensive to collect.

3D-aware GANs. Leveraging the recent developments of neural implicit representations [8, 30, 37, 44] and radiance fields [2, 24, 31, 32, 45], 3D GANs [1, 3–5, 9, 10, 15, 34, 35, 42, 46, 48, 52, 55–58] train generative 3D radiance fields from a set of single-view images. These methods render the sampled scenes from various viewpoints via volume rendering and supervise adversarially. Many of these approaches

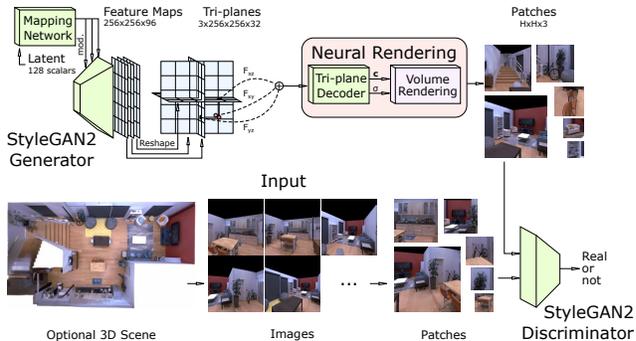


Figure 2. SinGRAF pipeline. The framework takes as input a few images of a single scene, for example rendered from a 3D scan or photographed (bottom). The 3D-aware generator (top left) is then trained to generate 2D feature planes that are arranged in a triplane configuration and rendered into patches of varying scale (top right). These rendered patches along with patches cropped from the input images are then compared by a discriminator. Once trained, the SinGRAF generator synthesizes different realizations of the 3D scene that resemble the appearance of the training images while varying the layout.

apply their discriminators on full-resolution images during training, but some also employ patch-based discriminators [28, 40, 46]. The resulting 3D GANs can create diverse 3D radiance fields that enable view-consistent NVS. Existing 3D GANs, however, rely on a large amount of training data, while our model only uses a few images of a single 3D scene.

Few-Shot Generative Models. Recently, researchers have started applying generative modeling techniques to few-shot settings, where only a few or single examples are given. In the 2D image domain, SinGAN and its extensions [19, 41, 54] explored the idea of training a CNN-based hierarchical generator on a single image, supervised using patch discrimination at multiple scales. The strategy of learning the internal patch distribution of a single example to train a generative model has been widely adopted for various tasks, including the synthesis of videos [16], motion sequences [23], 3D textures [17, 39], or 3D shapes [18, 51]. However, none of these works applied 3D generative models from single-scene images. Concurrently to our work, [50] trains generative radiance fields from single-scene images but focuses mainly on stochastic and repetitive synthetic scenes, rather than structured, human-made scenes.

3. Single Scene 3D GAN

Our system takes as input an unposed set of images taken from a single scene and outputs a 3D generative model G that can generate diverse 3D radiance fields. We assume the

intrinsic parameters of the camera are available. We do not assume our target scene to be static, i.e., we allow temporal changes in the scene. In the experiment section, we show our method’s behavior for dynamic scenes. We illustrate our image generation and discrimination processes in Fig. 2. Source code and pre-trained models will be made available.

3.1. Rendering Model

Our generative model $\mathbf{G}(\Gamma, \mathbf{z})$ takes a set of query rays Γ and a noise vector \mathbf{z} and outputs RGB predictions for the rays. Below we briefly discuss the rendering process of \mathbf{G} .

Generator Architecture \mathbf{G} represents continuous radiance fields using tri-planes, following [4]. We adopt a StyleGAN2-based generator backbone [22], which consists of a mapping network that takes as input a noise vector $\mathbf{z} \sim \mathbb{R}^{128}$ and transforms it into a latent code vector $\mathbf{w} \sim \mathbb{R}^{128}$. Next, a synthesis network transforms \mathbf{w} into a 2D feature image $\mathbf{F} \in \mathbb{R}^{N \times N \times 3C}$ with a total of $3C$ feature channels. Following [4], we split these feature channels into three axis-aligned feature planes $\mathbf{F}_{xy}, \mathbf{F}_{xz}, \mathbf{F}_{yz} \in \mathbb{R}^{N \times N \times C}$.

A color \mathbf{c} and density σ value can now be queried at an arbitrary 3D coordinate \mathbf{x} by aggregating the triplane features and processing them by a small multilayer perceptron-style decoder, $\text{MLP} : \mathbb{R}^{3C} \rightarrow \mathbb{R}^4$, as

$$(\mathbf{c}(\mathbf{x}), \sigma(\mathbf{x})) = \text{MLP}(\mathbf{F}_{xy}(\mathbf{x}) + \mathbf{F}_{xz}(\mathbf{x}) + \mathbf{F}_{yz}(\mathbf{x})). \quad (1)$$

Note that we do not model view-dependent effects.

Neural Rendering. Using volume rendering [26, 31], we project the 3D neural field into 2D images. For this purpose, the aggregated color $\mathbf{C}(\mathbf{r})$ of a ray $\mathbf{r} \in \Gamma$ is computed by integrating the field as

$$\mathbf{C}(\mathbf{r}) = \int_{t_n}^{t_f} T(t) \sigma(\mathbf{r}(t)) \mathbf{c}(\mathbf{r}(t)) dt, \quad (2)$$

$$T(t) = \exp\left(-\int_{t_n}^t \sigma(\mathbf{r}(s)) ds\right), \quad (3)$$

where t_n and t_f indicate near and far bounds along the ray $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$ pointing from its origin \mathbf{o} into direction \mathbf{d} . The continuous volume rendering equation (Eq. (3)) is typically computed using the quadrature rule [26].

Our volume rendering step directly outputs RGB color images or patches; we do not apply a superresolution module on the rendered values. Moreover, we use 96 samples per ray and do not use hierarchical ray sampling [31].

3.2. Training Process

Progressive Patch Scaling. Given a set of input images, most existing 3D GANs generate and discriminate images at full resolution. However, as shown in the ablation study



Figure 3. Visualizations of patches at different scales.

(Sec. 4.8), adversarial training at full scale leads to a collapse of the learned distribution to a single mode, likely because the joint information of the images uniquely determines a single 3D structure. Therefore, we turn to learn the internal patch distribution of the images, inspired by some existing works [41, 51, 54].

Directly extending hierarchical patch-based GANs [41, 51, 54], however, would require expensive training of a pyramid of generators with progressively-growing feature plane resolutions. To address the issue, we notice that our tri-plane features define a continuous radiance field and thus are able to render patches at an arbitrary resolution and scale. Therefore, we use a single generator network and continuously control the scale of the patches during training, forgoing the progressive training of multiple generator networks [41, 51, 54].

Given input images $\{I_1, \dots, I_N | I_i \in \mathbb{R}^{H' \times H' \times 3}\}$ of resolution $H' \times H'$ and field of view (FOV) θ , we consider the image plane \mathcal{P} of a virtual camera with the same FOV. During training, we volume-render patches with *fixed* resolution $H \times H$ but with varying scale s . Here, the scale $s \in [0, 1]$ indicates the spatial extent of a patch on the image plane of \mathcal{P} . When $s = 1.0$ the patch covers the entire image plane with $H \times H$ pixels, which is used for full-resolution training of existing 3D GANs [5, 40]. With smaller s , the patch will cover a smaller window in the image plane with the same resolution, thus containing more details. The location of the patch window on \mathcal{P} is sampled randomly. We render a sampled patch ρ with its associated rays $\Gamma_\rho: \mathbf{G}(\Gamma_\rho, \mathbf{z})$. Similarly, a ground truth patch can be sampled by cropping I_i 's with a given patch scale s (see Fig. 3).

The scale value s is sampled randomly for each patch from a uniform distribution $s \sim \mathcal{U}(s_{\min}(t), s_{\max}(t))$, where t is the current training epoch. We schedule the scale distribution so that in early epochs, we have larger patches to provide scene structural information and gradually decrease the scales towards the end of the training to induce better quality and diversity. Refer to supplementary for details.

Our discriminator network \mathbf{D} , whose architecture closely follows that of [22], takes as input patches and outputs a scalar, indicating the realism of the patches. Because our patches have varying scales, we additionally condition \mathbf{D} with the scale value s of each patch ρ :

$$\mathbf{D} : \mathbf{G}(\Gamma_\rho, \mathbf{z}) \times s \mapsto \mathbb{R}. \quad (4)$$

Implementation-wise, we simply repeat the scale value to match the patch resolution.

Data augmentation. Our approach aims at generating a 3D scene from a limited set of 2D observations. As such, it is desirable to augment the available data during training. To this end, we apply data augmentation techniques for both image and camera pose data.

In addition to the usual image augmentation techniques used for 2D GANs, such as translation, cropping, and cutout, we suggest a perspective augmentation approach, which is appropriate for generating 3D scenes with patch discrimination. Based on the known camera pose, we reproject image patches so as to imitate camera rotations followed by perspective projections. Camera rotations without changing the position do not induce occlusions or parallax; thus, this approach is applicable to captured content. Moreover, because our model discriminates patches instead of full images, unknown regions after rotation and perspective projection can be cropped. In practice, we start our training without rotational augmentation but gradually introduce and increase this augmentation up to 15° for later training epochs where the patch scale s decreases. Additional details on data augmentation are included in the supplement.

Camera Pose Distribution. To render patches from radiance fields we need to sample a virtual camera to render from. We define the camera pose distribution non-parametrically using a set of 1,000 cameras: $\mathcal{T} = \{T_1, \dots, T_{1000} | T_i \in SE(3)\}$; each virtual camera τ is randomly chosen from \mathcal{T} during training. We reject the sampled τ when the occupancy (opacity) value at the camera center is above some threshold. The T_i 's are initially sampled from a zero-mean 2D Gaussian on a plane with shared heights with random rotation about the vertical axis. During training, we jitter the T_i 's on translation and rotation.

The randomly initialized camera distribution, however, may not be a good representation of the real camera distribution of the input images. While our method can generate reasonable 3D scenes when training with random distributions, we find it beneficial to optimize the camera distribution \mathcal{T} for higher-quality outputs. Such optimization will transform the distribution to accommodate poses with various rotations. Therefore, we optimize the T_i 's via automatic differentiation using our adversarial loss in the early stage of the training. We eventually stop the optimization since smaller patches at the later stage contain weak pose information. In practice, we decompose the poses into more easily optimizable forms, which we discuss further in the supplementary.

Training Objective Besides the regular min-max adversarial loss [14], we find it useful to adopt two regularization

losses: R1 gradient penalty [29] and discriminator reconstruction loss [25] (\mathcal{E}_R). Combining the regularizations with the adversarial loss, we have the final objective as follows:

$$\mathcal{E}(\mathbf{G}, \mathbf{D}, \mathcal{T}) = \mathbb{E}_{\mathbf{z}, \tau \sim \mathcal{T}, s, \rho} [f(\mathbf{D}(\mathbf{G}(\Gamma_\rho, \mathbf{z})))] + \mathbb{E}_{I, s, \rho} [f(-\mathbf{D}(\rho, s)) + \lambda_1 |\nabla_\rho \mathbf{D}(\rho, s)|^2] + \lambda_2 \mathcal{E}_R(\mathbf{D}), \quad (5)$$

where $f(a) = -\log(1 + \exp(-a))$, and λ 's are balancing parameters. Note that the sampled patch ρ is a function of the sampled scale s . We minimize the objective using the ADAM optimizer with learning rate of $2e-3$.

4. Experiments

We conduct a number of experiments to test SinGRAF's capabilities to (i) create realistic 3D scene variations across a diverse set of challenging indoor scenes, (ii) reliably produce view-consistent 3D representations, and (iii) handle scene dynamics. At the end of the section, we show ablation studies and justify our design decisions. We refer to supplementary for additional empirical analysis and results.

4.1. Datasets

We test and compare our approach on five scenes from the Replica dataset and one scene from the Matterport3D dataset, featuring realistically scanned indoor scenes. Within the Replica dataset, we choose to cover a good range of different scene types, including offices, apartments, and hotel rooms. The Replica dataset, however, only contains scanned indoor scenes that are typical of their own categories. To stress-test our algorithm, we add a large 'castle' ballroom dataset from the Matterport3D dataset that has a very different appearance from 'regular' indoor scenes. For each of the above 6 scenes, we render 100 views by sampling camera locations from a zero-mean Gaussian distribution with random rotations about the height axis. We reject the cameras that collide with occupied volumes. Lastly, we showcase our method on a captured image dataset where we use our own photographs of an outdoor scene. We capture the images using a hand-held consumer-grade smartphone camera and pass them to our algorithm without intrinsic or extrinsic calibrations.

4.2. Baselines

As a baseline, we compare SinGRAF against the current state-of-the-art 3D scene generative method, GSN [10]. While numerous 3D-GAN methods were proposed, we could not find any other suitable baseline that demonstrated the modeling of apartment-scale scenes, such as the ones from the Replica dataset. Concurrently to our work, GAUDI [3] extended GSN to model even larger-scale scenes with generative modeling, but the code is not publicly available for us to run. To compare against GSN, we



Figure 4. Results of four difference scenes, as indicated. For each example, we show some of the input images, the result achieved by the GSN [10] baseline, and three realizations of SinGRAF’s results for the same input images. Note that GSN is mode-collapsed and not able to generate different realizations of this scene while SinGRAF is capable of generating a diverse set of realizations with high image quality.

	GSN (128 ²)		SinGRAF (128 ²)		SinGRAF (256 ²)	
	KID↓	Div.↑	KID↓	Div.↑	KID↓	Div.↑
office_3	.061	.001	.044	.297	.050	.378
hotel_0	.049	.012	.037	.413	.046	.490
apt.0	.069	.001	.037	.401	.049	.467
fri_aprt.4	.052	.001	.037	.335	.055	.408
castle	.050	.001	.064	.248	.088	.318
office_0	.075	.001	.053	.001	.062	.003
dynamic	.089	.013	.033	.298	.050	.365

Table 1. Quantitative Results. We measure the realism and diversity of the 3D scenes generated from SinGRAF and GSN on Replica and Matterport3D scenes. KID compares the distributional difference between the rendered and ground truth images of the scenes. We measure the diversity by rendering images with various latent vectors from fixed camera. Overall, we outperform the GSN baseline on both metrics in all but one case.

use their public codebase and run their algorithm using the same 100-image datasets. During training, we sample the virtual cameras from the 10,000 camera samples used in their original paper. Following the original implementation, the images are volume-rendered at 64×64 resolution and upsampled to 128×128 with a learned CNN.

4.3. Metrics

To measure the generated image quality we use the popular Kernel Inception Distance (KID) score, which reliably measures the distance between two sparsely sampled ($N \leq 500$) image distributions. This is in contrast with FID, which introduces significant biases in the low-data regime. We randomly sample 500 images for all methods in the same way as each of them samples during training. The 500 ground truth images are sampled as we generate the training images.

We measure the diversity of scene generation by sampling multiple images from a fixed camera and computing average LPIPS distance (average pair-wise LPIPS distances), following [20]. Note that, however, because our generated scenes are 3D, there exist loopholes to our diversity metric. For example, the diverse images generated from a fixed camera could be identical in the 3D space under rigid transformations. However, we did not observe this edge case and observe that the high diversity score from a fixed view leads to diversity beyond rigid transformations and vice versa.

4.4. Scene Generation Results

We showcase the visual quality and diversity of our trained generative model across scenes. As can be seen in Figs. 1, 4a, 4b, 4c, our method is able to synthesize plausible and realistic variants of the original scenes under a wide range of indoor scene environments. For example, the “office_3” scene shown in Fig. 1 contains a meeting table (or-

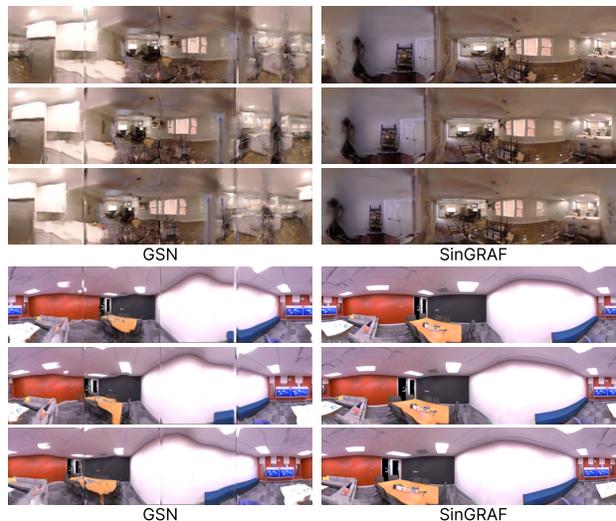


Figure 5. View consistency test with small rotations (15°). We check the 3D consistency of the neural fields generated by SinGRAF and GSN by visualizing the scene from varying viewpoints. GSN produces inconsistent discontinuities and abrupt structure changes (i.e., plant in the right), while SinGRAF demonstrates smooth and 3D-consistent view changes. Note that the shown panorama naturally induces some image distortions.

ange), a sofa (grey), and a coffee table (white). Note how our model is able to augment the room by duplicating, elongating, or rotating the meeting table and the sofa. For the “hotel_0” and “apartment_0” scenes shown in Figs. 4a, 4b, our model was able to capture *structural* diversity of the scenes while preserving the details and general appearance of the scene. As shown in the results of “fri_apartment_4” (Fig. 4c) our model was able to generate scenes with very different sizes by duplicating large scene structures. The results in the “castle” scene, shown in Fig. 4d, features various structural changes along with the diversity of the number and locations of the chairs in the scene. Overall, SinGRAF was able to synthesize a remarkable amount of variation across scenes, even when the scene is small and simple, e.g., “office_3” or “hotel_0.”

On the other hand, applying the strongest current baseline for scene generation, GSN [10], produces mode-collapsed results without diversity for all of the tested scenes. Quantitatively, as shown in Tab. 1, SinGRAF outperforms GSN both in terms of realism and diversity.

Video Results. We urge readers to watch our supplementary videos to fully appreciate the quality and diversity of our 3D scenes.

4.5. View-Consistent Scene Generation

We notice that even for the single-mode results of GSN, the scene renderings are of suboptimal quality, contain-

ing spurious blur (Fig. 4b) or unnaturally abrupt content changes by viewpoints (Fig. 4a). We hypothesize that GSN learned to generate plausible images but does not learn to create consistent 3D structures. To test this hypothesis, we visualize both GSN and our method when rotating the panorama twice by 15° each (Fig. 5). For the case of GSN, the slight change of viewpoints resulted in significant structural changes, for example changing the room or changing the shape of the furniture. In contrast, SinGRAF reliably generates consistent images and scene structures with varying viewpoints, indicating strong 3D awareness of the learned representations.

4.6. Modeling Scene Dynamics

We highlight that SinGRAF is especially robust in capturing scene dynamics since it does not rely on any pixel-wise reconstruction loss. To verify this claim, we run our method on the “frl_apartment” scene in the Replica dataset that has been captured with various configurations (5 different scene setups). The synthesis results shown in Fig. 6 indeed confirm SinGRAF’s ability to train a high-quality generative model on scenes that are not static. Note that, given diverse input configurations, SinGRAF was able to induce more variations to the scenes, by reorganizing the objects in the scenes. We note that training on these dynamic configurations did not result in diverse scene generation of GSN – it produced single mode outputs.

4.7. Towards Casually-captured Scenes

We also test our method “in the wild”, i.e., with captured content rather than pre-scanned scenes. For this experiment, we took 100 photographs of an outdoor apartment setting using a consumer-level smartphone. This scene is particularly challenging because it contains a lot of high-frequency textures, such as trees and grass, and a large dark window with view-dependent reflections. The camera setting also makes the problem difficult, as the intrinsic parameters are unknown and training images contain lens distortions. We approximate the field of view of the cameras with 65°, and our model successfully generates variations with visually pleasing quality, as shown in Fig. 7. The average LPIPS distance is 0.001 for GSN, and 0.372 and 0.444 for SinGRAF at 128² and 256² image resolution, respectively. Although a KID score is not available for this example, because we only have 100 training images and no ground truth, this experiment demonstrates the potential of SinGRAF to be applied in the wild.

4.8. Analysis

We analyze the effects of our design decisions through ablation studies next. The quantitative ablation results are displayed in Tab. 2.

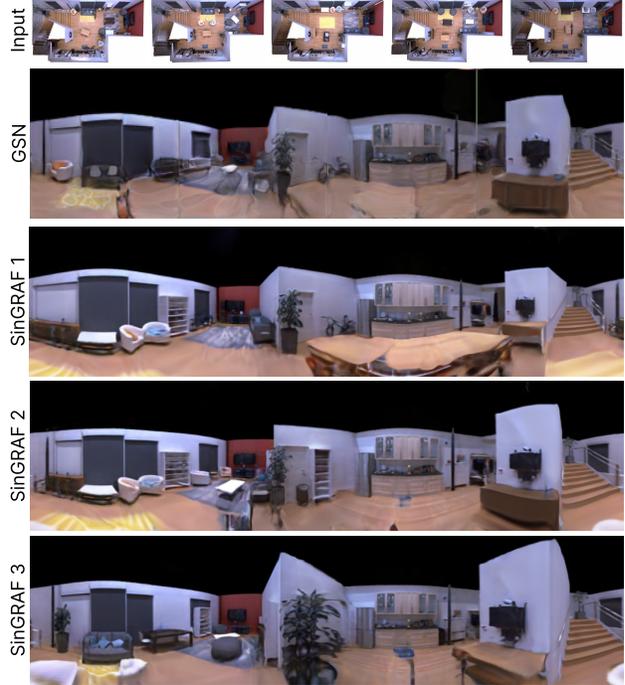


Figure 6. Modeling scene dynamics. We train our model using 500 images from five different scene configurations of the “frl_apartment”. Our resulting generative model produces 3D scenes with highly diverse variations of objects and structures, e.g., the locations of furniture.



Figure 7. Casually-captured scene. Given 100 input images photographed with a smartphone, our method is successful in generating high-quality variations of the 3D scene.

Patch Discrimination. As suggested by prior works on single-image GAN, learning the scene-internal patch statistics is crucial in inducing diversity into generated results. Indeed, without the patch discrimination, i.e., using full-scale ($s = 1.0$) images for adversarial training, results in a

	128 × 128		256 × 256	
	KID↓	Div.↑	KID↓	Div.↑
full & half-scale patches	.183	.001	NA	NA
progressive patches	.046	.308	.068	.374
+ camera opt.	.037	.295	.056	.368
+ perspective aug.	.037	.335	.055	.408

Table 2. Ablation study. We ablate our model to study the effects of patch discrimination, camera pose optimization, and perspective augmentation using the “f1r_lapartment_4” scene. Note that progressive patch scaling is essential for obtaining diverse scenes. Camera distribution optimization improves image quality while perspective augmentation maximizes scene diversity.

complete mode collapse of the model. Even when combining half-scale ($s = 0.5$) and full-scale ($s = 1.0$) patches as a 50% mix, the training results in no diversity (first row of Tab. 2), suggesting the importance of our progressive patch scaling strategy.

Camera Pose Optimization. We test the importance of our non-parametric camera distribution optimization scheme described in Sec. 3.2. As expected, adjusting the camera distributions result in higher realism of the generated outputs, resulting in lower KID scores (third row of Tab. 2). On the other hand, the optimization scheme only slightly hurts the diversity of the generation.

Perspective Augmentation. Given our low-data regime, a possible way of maximizing the diversity is via data augmentation strategies, e.g., perspective image augmentation (see Sec. 3.2). As expected, applying perspective image augmentation leads to higher diversity (fourth row of Tab. 2). What we found interesting is that adding perspective augmentation of 15° did not result in lower KID scores. This shows that perspective augmentation is an effective strategy in our setting.

Failure cases. We notice that for the “office_0” scene, shown in Fig. 8, SinGRAF mode-collapses and fails to generate diversity. We believe that this occurs because the detailed paintings on the walls uniquely identify the location of the patches in relation to others.



Figure 8. Failure case of the “office_0” scene. The lack of diversity for this scene is likely due to the paintings on the walls that uniquely determine the relative locations of most patches.

5. Discussion

In this work, we take first steps towards a 3D generative model for a single scene from a set of unposed images. Our model, SinGRAF, is able to generate diverse and realistic variations of the given scene while preserving the semantics of human-made structures. In contrast to 2D single-image GAN methods [41], which typically train a pyramid of generators, SinGRAF is trained with a single generator architecture, and is thus simple to implement and train. We achieve this by continuously adjusting the patch scales during training, leveraging the continuous neural fields representation. Given the unposed and scarce nature of our image data, we find it useful to optimize the camera pose distributions and apply perspective augmentations to the images.

Below we discuss three interesting implications of our approach to the field of 3D computer vision.

Reconstructing Variations. SinGRAF, in a sense, reconstructs the distribution of plausible 3D scenes given a set of images. Traditionally, the 3D vision community largely focused on finding the *most likely* mode of the 3D scene for reconstruction. Relaxing this constraint to instead sample from the posterior distribution of scenes could lead to new technologies or applications.

Unposed Reconstruction. SinGRAF does not use estimated poses of the input images as it relies on an adversarial rather than a pixel-wise loss. While SinGRAF, in its current form, is not designed to accurately reproduce the ground truth 3D scene, it might be extended to provide control of the *narrowness* of the distribution, depending on the applications.

Dynamic Scene Reconstruction. In contrast to traditional 3D reconstruction methods, scene dynamics would likely *improve* the quality and diversity of SinGRAF training (see Fig. 6). Such a trend implies that it could be easier to extend SinGRAF to operate on highly dynamics scenes, e.g., rock concerts, rather than developing physical models to handle complex scene dynamics.

Limitations. The quality and diversity of the SinGRAF outputs depend on the input images, scenes, and choice of views, which are hard to predict or control. Moreover, SinGRAF training is currently expensive and takes 1–2 days per scene with a single RTX 6000 GPU, although it is comparable to existing 3D GANs (e.g., GSN). We expect improvements in the training speed with the development of more efficient continuous representations, such as [7, 32].

Acknowledgements

This project was in part supported by Samsung, Stanford HAI, a PECASE from the ARO, ARL grant W911NF-21-2-0104, a Vannevar Bush Faculty Fellowship, and gifts from the Adobe and Snap Corporations.

References

- [1] Sherwin Bahmani, Jeong Joon Park, Despoina Paschalidou, Hao Tang, Gordon Wetzstein, Leonidas Guibas, Luc Van Gool, and Radu Timofte. 3D-aware video generation. *arXiv preprint arXiv:2206.14797*, 2022. [2](#)
- [2] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-NeRF: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5855–5864, 2021. [2](#)
- [3] Miguel Angel Bautista, Pengsheng Guo, Samira Abnar, Walter Talbott, Alexander Toshev, Zhuoyuan Chen, Laurent Dinh, Shuangfei Zhai, Hanlin Goh, Daniel Ulbricht, Afshin Dehghan, and Josh Susskind. GAUDI: A neural architect for immersive 3D scene generation. *arXiv*, 2022. [2](#), [4](#)
- [4] 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. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2022. [2](#), [3](#)
- [5] Eric R Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. pi-GAN: Periodic implicit generative adversarial networks for 3D-aware image synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5799–5809, 2021. [1](#), [2](#), [3](#)
- [6] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Habber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3D: Learning from RGB-D data in indoor environments. *International Conference on 3D Vision (3DV)*, 2017. [2](#)
- [7] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. TensorRF: Tensorial radiance fields. *arXiv preprint arXiv:2203.09517*, 2022. [8](#)
- [8] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5939–5948, 2019. [2](#)
- [9] Yu Deng, Jiaolong Yang, Jianfeng Xiang, and Xin Tong. GRAM: Generative radiance manifolds for 3D-aware image generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2022. [2](#)
- [10] Terrance DeVries, Miguel Angel Bautista, Nitish Srivastava, Graham W. Taylor, and Joshua M. Susskind. Unconstrained scene generation with locally conditioned radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14304–14313, October 2021. [2](#), [4](#), [5](#), [6](#)
- [11] David S Ebert, F Kenton Musgrave, Darwyn Peachey, Ken Perlin, and Steven Worley. *Texturing & modeling: a procedural approach*. Morgan Kaufmann, 2003. [1](#), [2](#)
- [12] Matthew Fisher, Daniel Ritchie, Manolis Savva, Thomas Funkhouser, and Pat Hanrahan. Example-based synthesis of 3D object arrangements. *ACM Transactions on Graphics (TOG)*, 31(6):1–11, 2012. [2](#)
- [13] Thomas Funkhouser, Michael Kazhdan, Philip Shilane, Patrick Min, William Kiefer, Ayellet Tal, Szymon Rusinkiewicz, and David Dobkin. Modeling by example. *ACM transactions on graphics (TOG)*, 23(3):652–663, 2004. [1](#), [2](#)
- [14] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020. [2](#), [4](#)
- [15] Jiatao Gu, Lingjie Liu, Peng Wang, and Christian Theobalt. StyleNeRF: A style-based 3D-aware generator for high-resolution image synthesis. *arXiv preprint arXiv:2110.08985*, 2021. [2](#)
- [16] Niv Haim, Ben Feinstein, Niv Granot, Assaf Shocher, Shai Bagon, Tali Dekel, and Michal Irani. Diverse generation from a single video made possible. In *European Conference on Computer Vision*, pages 491–509. Springer, 2022. [2](#)
- [17] Philipp Henzler, Niloy J Mitra, and Tobias Ritschel. Learning a neural 3D texture space from 2D exemplars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8356–8364, 2020. [2](#)
- [18] Amir Hertz, Rana Hanocka, Raja Giryes, and Daniel Cohen-Or. Deep geometric texture synthesis. *arXiv preprint arXiv:2007.00074*, 2020. [2](#)
- [19] 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. [2](#)
- [20] Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. Multimodal unsupervised image-to-image translation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018. [6](#)
- [21] Evangelos Kalogerakis, Siddhartha Chaudhuri, Daphne Koller, and Vladlen Koltun. A probabilistic model for component-based shape synthesis. *Acm Transactions on Graphics (TOG)*, 31(4):1–11, 2012. [1](#), [2](#)
- [22] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8110–8119, 2020. [2](#), [3](#)
- [23] Peizhuo Li, Kfir Aberman, Zihan Zhang, Rana Hanocka, and Olga Sorkine-Hornung. GANimator: Neural motion synthesis from a single sequence. *arXiv preprint arXiv:2205.02625*, 2022. [2](#)
- [24] David B Lindell, Dave Van Veen, Jeong Joon Park, and Gordon Wetzstein. BACON: Band-limited coordinate networks for multiscale scene representation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16252–16262, 2022. [2](#)
- [25] Bingchen Liu, Yizhe Zhu, Kunpeng Song, and Ahmed Elgammal. Towards faster and stabilized GAN training for high-fidelity few-shot image synthesis. In *International Conference on Learning Representations*, 2020. [4](#)
- [26] N. Max. Optical models for direct volume rendering. *TVCG*, 1995. [3](#)

- [27] Radomír Měch and Przemyslaw Prusinkiewicz. Visual models of plants interacting with their environment. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 397–410, 1996. 1, 2
- [28] Quan Meng, Anpei Chen, Haimin Luo, Minye Wu, Hao Su, Lan Xu, Xuming He, and Jingyi Yu. GNeRF: GAN-based neural radiance field without posed camera. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6351–6361, 2021. 2
- [29] Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for GANs do actually converge? In *International conference on machine learning*, pages 3481–3490. PMLR, 2018. 4
- [30] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3D reconstruction in function space. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4460–4470, 2019. 2
- [31] 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. *Communications of the ACM*, 65(1):99–106, 2021. 1, 2, 3
- [32] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (TOG)*, 41(4):1–15, 2022. 2, 8
- [33] F Kenton Musgrave, Craig E Kolb, and Robert S Mace. The synthesis and rendering of eroded fractal terrains. *ACM SIGGRAPH Computer Graphics*, 23(3):41–50, 1989. 2
- [34] Michael Niemeyer and Andreas Geiger. GIRAFFE: Representing scenes as compositional generative neural feature fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11453–11464, 2021. 2
- [35] Roy Or-El, Xuan Luo, Mengyi Shan, Eli Shechtman, Jeong Joon Park, and Ira Kemelmacher-Shlizerman. StyleSDF: High-resolution 3D-consistent image and geometry generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13503–13513, 2022. 2
- [36] Yoav IH Parish and Pascal Müller. Procedural modeling of cities. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 301–308, 2001. 2
- [37] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 165–174, 2019. 2
- [38] Despoina Paschalidou, Amlan Kar, Maria Shugrina, Karsten Kreis, Andreas Geiger, and Sanja Fidler. ATISS: Autoregressive transformers for indoor scene synthesis. *Advances in Neural Information Processing Systems*, 34:12013–12026, 2021. 2
- [39] Tiziano Portenier, Siavash Arjomand Bigdeli, and Orcun Goksel. GramGAN: Deep 3D texture synthesis from 2D exemplars. *Advances in Neural Information Processing Systems*, 33:6994–7004, 2020. 2
- [40] Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. GRAF: Generative radiance fields for 3D-aware image synthesis. *Advances in Neural Information Processing Systems*, 33:20154–20166, 2020. 1, 2, 3
- [41] Tamar Rott Shaham, Tali Dekel, and Tomer Michaeli. SinGAN: Learning a generative model from a single natural image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. 2, 3, 8
- [42] Zifan Shi, Yujun Shen, Jiapeng Zhu, Dit-Yan Yeung, and Qifeng Chen. 3D-aware indoor scene synthesis with depth priors. In *ECCV*, 2022. 2
- [43] J. R. Shue, E. R. Chan, R. Po, Z. Ankner, J. Wu, and G. Wetzstein. 3D neural field generation using triplane diffusion. In *CVPR*, 2023. 2
- [44] Vincent Sitzmann, Julien N.P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In *Proc. NeurIPS*, 2020. 2
- [45] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous 3D-structure-aware neural scene representations. In *Advances in Neural Information Processing Systems*, 2019. 2
- [46] Ivan Skorokhodov, Sergey Tulyakov, Yiqun Wang, and Peter Wonka. EpiGRAF: Rethinking training of 3D GANs. *arXiv preprint arXiv:2206.10535*, 2022. 2
- [47] Julian Straub, Thomas Whelan, Lingni Ma, Yufan Chen, Erik Wijmans, Simon Green, Jakob J Engel, Raul Mur-Artal, Carl Ren, Shobhit Verma, et al. The Replica dataset: A digital replica of indoor spaces. *arXiv preprint arXiv:1906.05797*, 2019. 2
- [48] Ayush Tewari, Mallikarjun B R, Xingang Pan, Ohad Fried, Maneesh Agrawala, and Christian Theobalt. Disentangled3D: Learning a 3D generative model with disentangled geometry and appearance from monocular images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1516–1525, June 2022. 2
- [49] 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. 2
- [50] Yujie Wang, Xuelin Chen, and Baoquan Chen. SinGRAV: Learning a generative radiance volume from a single natural scene. *arXiv preprint arXiv:2210.01202*, 2022. 2
- [51] Rundi Wu and Changxi Zheng. Learning to generate 3D shapes from a single example. *arXiv preprint arXiv:2208.02946*, 2022. 1, 2, 3
- [52] Jianfeng Xiang, Jiaolong Yang, Yu Deng, and Xin Tong. GRAM-HD: 3D-consistent image generation at high resolution with generative radiance manifolds, 2022. 2
- [53] Kai Xu, Hao Zhang, Daniel Cohen-Or, and Baoquan Chen. Fit and diverse: Set evolution for inspiring 3D shape galleries. *ACM Transactions on Graphics (TOG)*, 31(4):1–10, 2012. 2
- [54] Rui Xu, Xintao Wang, Kai Chen, Bolei Zhou, and Chen Change Loy. Positional encoding as spatial inductive bias in GANs. In *Proceedings of the IEEE/CVF Conference*

- on *Computer Vision and Pattern Recognition (CVPR)*, pages 13569–13578, June 2021. [2](#), [3](#)
- [55] Yinghao Xu, Sida Peng, Ceyuan Yang, Yujun Shen, and Bolei Zhou. 3D-aware image synthesis via learning structural and textural representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 18430–18439, 2022. [2](#)
- [56] Yang Xue, Yuheng Li, Krishna Kumar Singh, and Yong Jae Lee. GIRAFFE HD: A high-resolution 3D-aware generative model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2022. [2](#)
- [57] Xuanmeng Zhang, Zhedong Zheng, Daiheng Gao, Bang Zhang, Pan Pan, and Yi Yang. Multi-view consistent generative adversarial networks for 3D-aware image synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2022. [2](#)
- [58] Peng Zhou, Lingxi Xie, Bingbing Ni, and Qi Tian. CIPS-3D: A 3D-aware generator of GANs based on conditionally-independent pixel synthesis. *arXiv preprint arXiv:2110.09788*, 2021. [2](#)