

# ClimateNeRF: Extreme Weather Synthesis in Neural Radiance Field

Yuan Li<sup>1,2\*</sup> Zhi-Hao Lin<sup>1\*</sup> David Forsyth<sup>1</sup> Jia-Bin Huang<sup>3</sup> Shenlong Wang<sup>1</sup>  
<sup>1</sup>University of Illinois Urbana-Champaign <sup>2</sup>Zhejiang University  
<sup>3</sup>University of Maryland, College Park

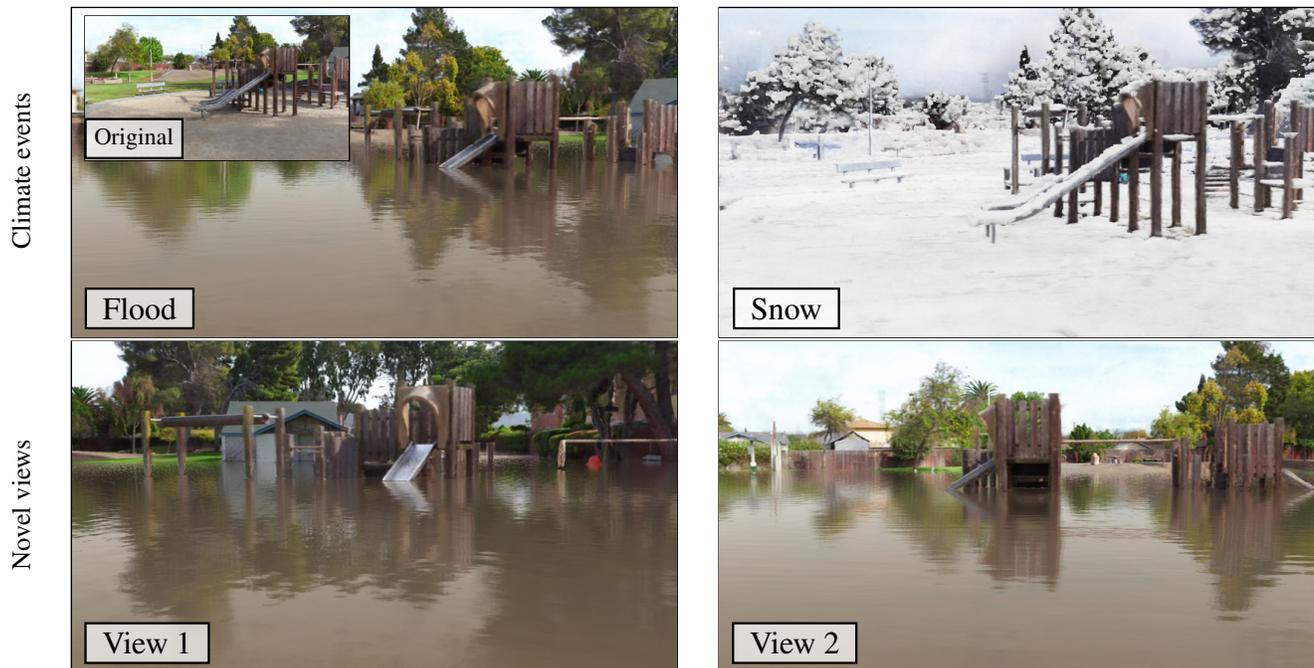


Figure 1: **Extreme weather synthesis.** We fuse NeRF modeling and physical simulation to produce 3D consistent renderings of scenes with simulated physical effects. We apply our method to climate effect simulation: what will it look like if the playground floods? or is covered in snow? Note in particular reflections and ripple effects on water; accumulated snow on horizontal surfaces; trees darkened by wintertime; and consistency of geometry (but not ripples!). [Project page](#)

## Abstract

*Physical simulations produce excellent predictions of weather effects. Neural radiance fields produce SOTA scene models. We describe a novel NeRF-editing procedure that can fuse physical simulations with NeRF models of scenes, producing realistic movies of physical phenomena in those scenes. Our application – Climate NeRF – allows people to visualize what climate change outcomes will do to them.*

*ClimateNeRF allows us to render realistic weather effects, including smog, snow, and flood. Results can be controlled with physically meaningful variables like water level. Qualitative and quantitative studies show that our simulated results are significantly more realistic than those from SOTA 2D image editing and SOTA 3D NeRF stylization.*

\* Both authors contributed equally to this research.

## 1. Introduction

This paper describes a novel procedure that fuses graphic simulations with NeRF models [50, 51] of scenes to produce realistic movies of physical phenomena in those scenes. We apply our method to produce compelling simulations of possible extreme weather outcomes – what would the playground look like after a minor flood? a major flood? a blizzard?

Our application is aimed at an important problem. Cumulative small changes are hard to reason about, and most people find it difficult to visualize what climate change outcomes will do to them [9, 23, 26, 37]. Steps to slow CO<sub>2</sub> emissions (say, reducing fossil fuel use) or to moderate outcomes (say, building flood control measures) come with immediate costs and distant benefits. It is hard to support such steps if one can’t visualize their effects.

Traditional graphic simulations can produce realistic

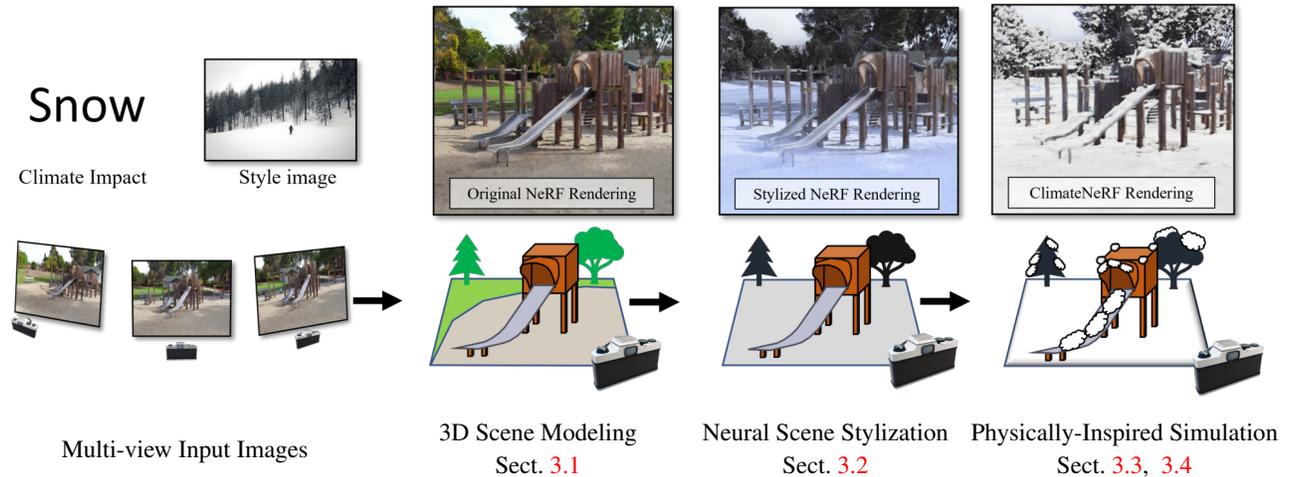


Figure 2: **Method overview.** Our method takes multiple posed images, the targeted climate event simulation (e.g., snow), and optionally a user-selected style image as inputs. First, we reconstruct the 3D scene using instant NeRF [51] (a variant of NeRF) (Sect. 3.1). The reconstructed radiance fields allow us to synthesize high-quality novel view imagery of the scene efficiently. Second, we optionally finetune the learned instant-NeRF model so that it captures the styles of the provided style image (Sect. 3.2). Such 3D consistent stylization is particularly useful for modeling weather effects that are hard to capture via physical simulation. Third, we simulate the climate events by integrating the relevant physical entities (snow, water, smog) to the scene and rendering physically plausible images.

weather effects for 3D scenes in a conventional simulation pipeline [62, 19, 27, 85, 15, 17, 24]. But these methods operate on conventional polygon models. Building polygon models that produce compelling renderings from a few images of a scene remains challenging. Neural radiance fields (NeRFs) produce photorealistic 3D scene models from few images [50, 4, 51, 10, 64, 72, 41] but (as far as we know) have rarely been investigated together with graphics simulations. Our method draws from an extensive literature, reviewed below, that explores editing these models.

Although there are numerous variants of NeRF, most are primarily focused on the task of novel-view synthesis. Integrating existing NeRF models for graphics simulation yields unsatisfactory results (Fig. 4). Specifically, we found the geometry of current NeRF models is not sufficiently accurate to enable physics-inspired interactions, and they cannot provide a realistic and complete illumination field to relight the inserted physics entities, such as water and snow. To address this challenge, we have systematically investigated various techniques to improve NeRF models, focusing on answering the critical research question of “*what constitutes a good simulatable NeRF representation?*”. Our resulting NeRF variant, ClimateNeRF, boasts high-quality geometry, a complete illumination field, and rich semantics, making it particularly well-suited for complex downstream editing and simulation tasks (Sec. 3.1).

ClimateNeRF allows us to render realistic weather effects, including smog, snow, and flood. These effects are consistent over frames so that compelling movies result. At a high level, we: adjust scene images to reflect the global

effects of physics; build a NeRF model of a scene from those adjusted images; recover an approximate geometric representation; apply the physical simulation in that geometry; then render using a novel ray tracer. Adjusting the images is essential. For example, trees tend to have less saturated images in winter. We use a novel style transfer approach in an NeRF framework to obtain these global effects without changing scene geometry (Sec. 3.2). Our ray tracer merges the physical and NeRF models by carefully accounting for ray effects during rendering (Sec. 3.3). An eye ray could, for example, first encounter a high NeRF density (and so return the usual result); or it could strike an inserted water surface (and so be reflected to query the model again).

We demonstrate the effectiveness of ClimateNeRF in various 3D scenes from the Tanks and Temple, MipNeRF360, and KITTI-360 datasets [4, 33, 40]. We compared to the state-of-the-art 2D image editing methods, such as stable diffusion inpainting [55], ClimateGAN [60]; state-of-the-art 3D NeRF stylization [78]. Both qualitative and quantitative studies show that our simulated results are significantly more realistic than the other competing methods. Furthermore, we also demonstrate the controllability of our physically-inspired approaches, such as changing the water level, wind strength and direction, and thickness of snow and smog.

Our approach results in view consistency (so we can make movies, which is difficult to do with frame-by-frame synthesis); compelling photorealism (because the scene is a NeRF representation); and is controllable (because we can adjust physically meaningful parameters in the simulation). As Fig. 1 illustrates, results are photo-realistic, physically

plausible, and temporally consistent.

Our contributions are three-fold:

1. We investigate the challenging problem of repurposing NeRF for realistic simulation and develop a novel framework that can effectively reduce NeRF modeling errors and facilitate graphical simulation.
2. We propose ClimateNeRF, a novel solution that can render realistic weather effects, including smog, snow, and flood, over real-world images.
3. We demonstrate that our proposed approach can produce compelling, photorealistic, geometrically consistent, and controllable free-viewpoint movies.

## 2. Related Work

**Climate simulation:** The importance of making climate simulations accessible is well known. [59] collects climate images dataset and performs image editing with CycleGAN [83]. [16, 60] leverage depth information to estimate water mask and perform GAN-based image editing and inpainting. [25] simulate fog and snow. These methods offer realistic effects for a single image but cannot provide immersive, view-consistent climate simulation. In contrast, ClimateGAN allows the view to move without artifacts.

**Novel view synthesis:** Neural Radiance Field (NeRF) [50, 3, 4, 68] leverage differentiable rendering for scene reconstruction producing photorealistic novel view synthesis. Training and rendering can be accelerated using multiple smaller MLPs and customized CUDA kernel functions [54, 41]. Sample efficiency can be improved by modeling the light field [2] or calculating ray intersections with explicit geometry, such as voxel grids [44, 72, 54], octrees [76], planes [71, 41], and point cloud [74, 86]. Explicit geometric representations improve efficiencies [63, 22, 47, 74, 80]. We build on Instant-NGP [51] as it offers fast learning and rendering, is more memory efficient than grid-like structures, and accelerates our simulation pipeline.

**Manipulating neural radiance fields:** NeRF entangles lighting, geometry, and materials, making appearance editing hard. A line of work [61, 56, 81, 79, 39, 7, 8] decomposes neural scene representations into familiar scene variables (like environment lighting, surface normal, diffuse color, or BRDF). Each can then be edited with predictable results on rendering. Image segmentation and inpainting can be “lifted” to 3D, allowing object removal [35, 6]. An alternative is to edit source images with text [70]. In contrast, our method does not directly edit the NeRF while producing substantial changes to scene appearance. Geometry editing by transforming a NeRF into a mesh and manipulating that (so editing shapes and compositing) has been demonstrated [75, 77, 14]. In contrast, our method renders the original NeRF together with simulation results.

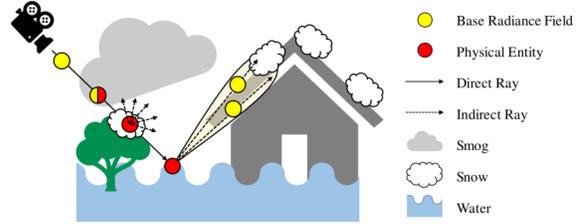


Figure 3: **Rendering Procedure of ClimateNeRF.** We first determine the position of physical entities (smog particle, snow balls, water surface) with physical simulation. We can then render the scene with desired effects by modeling the light transport between the physical entities and the scene. More specifically, we follow the volume rendering process and fuse the estimated color and density from 1) the original radiance field (by querying the trained instant-NGP model) and 2) the physical entities (by physically based rendering). Our rendering procedure thus maintain the realism while achieving complex, yet physically plausible visual effects.

**Style transfer:** Data-driven 2D stylization [31, 53, 28] is a powerful tool to change image color and texture, and can even change the content following masks and labels provided by users [43, 42, 84, 1]. Diffusion models [57, 55, 58] can generate extraordinary image quality based on text input. However, these 2D-based methods cannot currently generate 3D-consistent view synthesis; ours can. Another line of work [13, 11, 78] performs stylization on neural radiance fields. Given a style image, one changes the scene color into the same style. This line of work focuses on artistic effects; in contrast, our changes are driven by physical simulations.

**Physically-based simulation of weather:** Physical simulation of weather in computer graphics has too long a history for a comprehensive review. Fournier and Reeves obtain excellent capillary wave simulations with simple Fourier transform reasoning [20] (and we adopt their method), with modifications by [27, 66]; [18] simulates smoke; and [52, 19, 62] simulate snow in the wind using metaballs and fluid dynamics. Our method demonstrates how to benefit from such simulations while retaining the excellent scene modeling properties of NeRF-style models.

## 3. Method

ClimateNeRF fuses physical simulations with NeRF models of scenes to produce realistic videos of climate change effects. A simple example illustrates how components interact in our approach. Imagine we wish to build a model of a flooded scene in the Fall. We acquire images, apply a Fall style (Section 3.2), and build a NeRF from the results (Section 3.1). We then use geometric information in that NeRF to compute a water surface. This is represented using a density field, a color field together with normal and BRDF representations (Section 3.3). Finally, to render, we query

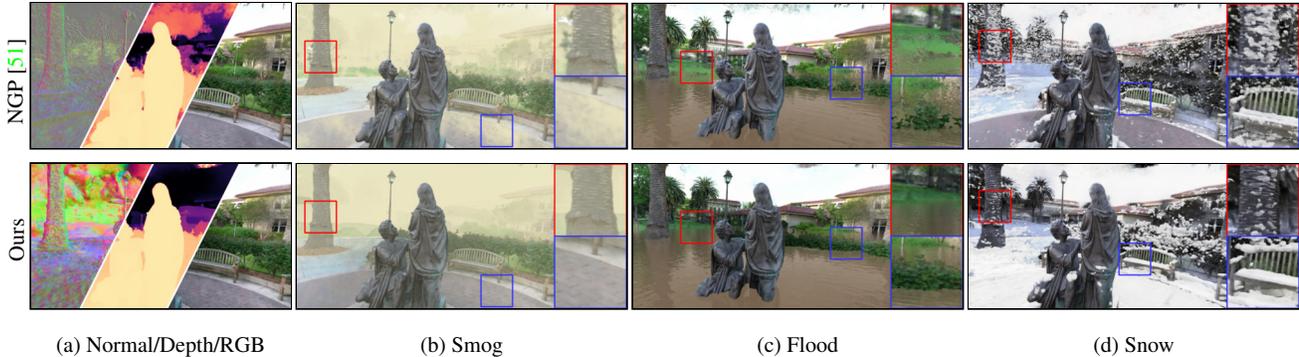


Figure 4: **Climate simulations reveal NeRF modeling errors.** Top: NGP, bottom: our NeRF model. On the left, there are normal, depth map, and RGB images. Although the rendered RGB images of the two NeRF versions appear similar, NGP exhibits depth errors in the sky regions and more normal noises on surfaces. This results in geometric inaccuracies in fog and flood images, as well as unrealistic snow accumulation.

the model with rays. The details are elaborate (Section 3.4), but the general idea is straightforward: we edit the NeRF’s density and color functions to represent effects like flood; and we intercept rays to represent specular effects. So if a ray encounters high density in the NeRF first, we use the NeRF integral for that ray; but if the first collision is with the water surface, we reflect that ray in the water surface, then query the NeRF with the reflected ray. Fig. 2 provides an overview of our approach.

### 3.1. Repurposing NeRF for Simulation

NeRF builds a parametric scene representation that supports realistic rendering from multiple images of a scene obtained at known poses. The scene is represented by a field  $(\sigma, \mathbf{c}) = F_{\theta}(\mathbf{x}, \mathbf{d})$ , which accepts position  $\mathbf{x}$  and direction  $\mathbf{d}$  and predicts density  $\sigma \in \mathbb{R}$  and color  $\mathbf{c} \in \mathbb{R}^3$ . This function is encoded in a multi-layer perceptron (MLP) with learnable parameters  $\theta$ . Rendering is done by querying radiance along appropriate choices of ray, computed as a volume integral [30]. This integral is estimated by drawing samples along the ray, evaluating the density and color of those samples, then accumulating values. The rendering process is differentiable so that NeRF can be trained by minimizing the image reconstruction loss over training views through gradient descent:  $\min_{\theta} \sum_{\mathbf{r}} \|C(\mathbf{r}) - C_{gt}(\mathbf{r})\|_2^2$ .

**Repurposing NeRF for simulation:** Although there are numerous variants of NeRF, most are primarily focused on the task of novel-view synthesis [50, 3, 4, 51, 10]. However, combining graphical simulation with these NeRF representations directly results in rendering artifacts, as shown in Fig. 4. These artifacts arise because the NeRF model cannot capture the scene’s accurate and clean geometry, resulting in floaters, and cannot complete the 4D light field, resulting in holes. When physical simulation, such as adding flood and snow into the scene, is incorporated, the artifacts are particularly

notable since human perception is more sensitive to these unnatural phenomena.

A good NeRF model for simulation must meet three requirements: 1) accuracy and completeness of geometry, which allows the scene to interact with physical objects; 2) semantic awareness, which enables the interaction to reflect the scene’s characteristics; and 3) coherence and completeness of the light field, which enables the use of NeRF for realistic lighting of the inserted physical entities. To this end, we systematically study multiple variants of techniques to improve NeRF. We use instant-NGP [51] as our base representation to reconstruct the scene for its efficiency and flexibility. This NeRF alternative explicitly stores multi-resolution features for scene representation, making it suitable as we can edit local features easily.

A physical simulation needs access to the surface normal of any point to compute interactions with snow and water, and it needs access to the point’s semantics (in the sense of semantic segmentation) to transfer style. We expand the NGP and allow it to output both semantic logic  $s$  and surface normal  $\mathbf{n}$ . There is no semantic or surface normal ground truth during training. We use an off-the-shelf pretrained monocular semantic segmentation network [73] to produce semantic maps for each image. We use density gradients  $\hat{\mathbf{n}} = -\frac{\nabla \sigma}{\|\nabla \sigma\|}$  [7, 61] (cf. Ref-NeRF [68]) to guide the predicted surface normals  $\mathbf{n}$  with a weighted MSE loss.

To simulate (say) a blizzard, we must add snow to the scene and turn trees dark, but we should not change the shape of the house. To keep spatial features intact at the stylization stage, we disentangle instant-NGP’s latent feature (as in [10, 63]). For each voxel in the NGP model, we split the latent feature into geometry features  $\gamma_{\text{geo}}$  and appearance features  $\gamma_{\text{app}}$ . The geometry features are trained to render density. The appearance features are used for rendering color, semantics, and normals. We will freeze the geometry feature vector during the stylization stage and change only

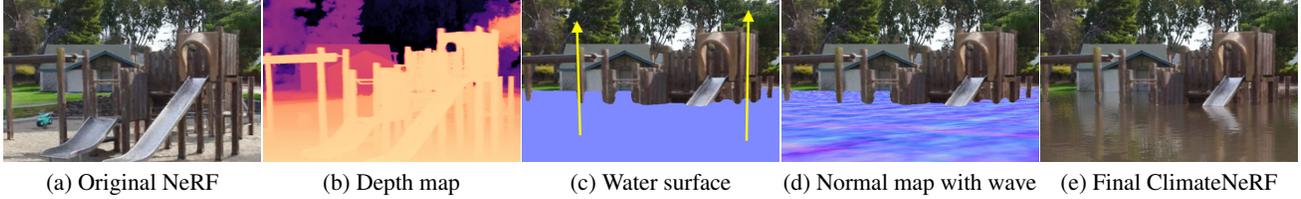


Figure 5: **Flooding simulation.** We first estimate the vanishing point direction based on the original image (a) and depth (b). With the vertical vanishing direction (yellow arrows painted (c)), we can insert a planar water surface  $\mathbf{n}_w(x - \mathbf{o}_w)$ . We use FFT based water surface simulation to produce a spatiotemporal surface normal map in (d). Our ClimateNeRF renders the scene with the simulated flood through ray tracing NeRF (e).

Method	Appearance Embed. [48]	Distortion Loss [4]	Predicted Normal	Predicted Semantics [82]	Sky Loss
NGP [51]					
Mip-NeRF360 [4]	✓	✓			
NeRFacto [65]	✓	✓	✓		
Ours	✓	✓	✓	✓	✓

Table 1: **Improving NeRF for simulation.**

the appearance feature vector.

### Improving NeRF geometry to reduce simulation errors:

The non-ideal geometry in NeRF, as seen in [50, 51], produces severe artifacts in simulation. To address this issue, we incorporate per-frame appearance embedding [48] and distortion loss [4] to eliminate floaters. Additionally, we utilize a surface normal MLP to enhance the smoothness of the scene normals while introducing a novel sky loss function  $\mathcal{L}_{\text{sky}} = \sum_{\mathbf{r} \in \text{sky}} e^{-d(\mathbf{r})}$ , where  $d(\mathbf{r})$  stands for depth along ray  $\mathbf{r}$  and  $\mathbf{r} \in \text{sky}$  corresponds to pixels labeled as sky, to eliminate floaters and fill in holes in the sky. Through these improvements in geometry, we have significantly improved the realism of the simulation (more details can be found in the supplementary material).

Our final NeRF renders density  $\sigma$ , color  $\mathbf{c}$ , semantic logit  $\mathbf{s}$  and normal  $\mathbf{n}$ , given a query point and ray direction:

$$(\sigma, \mathbf{c}, \mathbf{s}, \mathbf{n}) = F_\theta(\mathbf{x}, \mathbf{d}; \gamma_{\text{geo}}, \gamma_{\text{app}}). \quad (1)$$

Table. 1 compares the techniques used in NeRF variants and summarize what we find useful to improve the simulation.

### 3.2. Stylization

Deciduous trees drop their leaves in winter, and physical simulation is not efficient for capturing effects like this. We use FastPhotoStyle [38] to transfer style to rendered images from a pre-trained model  $F_\theta$ . We only transfer only to regions 'terrain', 'vegetation', or 'sky' regions to mimic natural weather change phenomena. The resulting images look realistic but are not necessarily view-consistent. Hence, a student instant-NGP model is fine-tuned to ensure the view consistency of the style-transferred scene. This is trained to minimize the color difference between our student NeRF-rendered results and style-transferred images. We keep the geometry intact and alter only appearance to achieve this goal, so only the appearance

feature code  $\gamma_{\text{app}}$  is optimized during the style transfer stage:  $\min_{\gamma_{\text{app}}} \sum_{\mathbf{r} \in \mathcal{R}} \|C(\mathbf{r}) - C_{\text{stylized}}(\mathbf{r})\|_2^2$ , where  $C(\mathbf{r})$  is render color and  $C_{\text{stylized}}$  is the style-transferred in the same view. This gives us a new NGP model  $(\sigma, \mathbf{c}') = F'_\theta(\mathbf{x}, \mathbf{d}, \ell_i^{(a)})$  which encodes the style. This optional style transfer step simulates composite effects, such as a flood in Fall, where the original images were captured in Spring.

### 3.3. Representing and Rendering Climate Effects

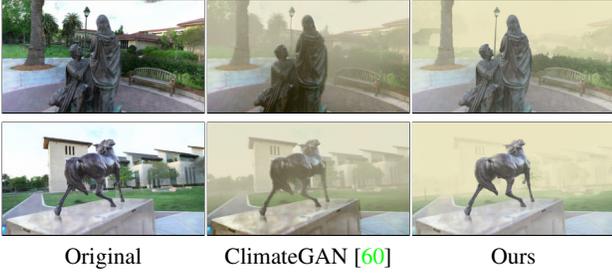
We want to generate scenes with new physical entities – snow, water, smog – in place. We must determine where they are (the job of physical simulation) and what the resulting image looks like (the job of rendering). *How* the simulation represents results is important because results must be accessible to the rendering process. Rendering will always involve computing responses to ray queries, so computing radiance at  $\mathbf{u}$  in direction  $\mathbf{v}$ . We must represent simulation results in terms of densities, and we must be able to compute normals and surface reflectance properties. Generally, we write  $O_\phi(\mathbf{x}; F_\theta) : \mathbb{R}^3 \rightarrow \mathbb{R}$  for a density resulting from a physical simulation;  $N_\phi(\mathbf{x}; F_\theta) : \mathbb{R}^3 \rightarrow \mathbb{S}^2$  for normals; and  $B_\phi(\mathbf{x}, \omega_o, \omega_i; F_\theta) : \mathbb{R}^9 \rightarrow \mathbb{R}$  for BRDF. Each depends on the existing scene  $F_\theta$ . Choice of  $B_\phi$  can simulate various effects, including the atmospheric effect of smog, refraction and reflections on water surfaces, and scattering of accumulated snow.  $\{O_\phi, N_\phi, B_\phi\}$  differs drastically across different physical simulations (details per effect in Sect. 3.4).

Once the physical entities are defined by functions  $O_\phi, N_\phi, B_\phi$ , we can render them into the image by modeling the light transport between the physical entities and the scene. Given the query point position  $\mathbf{x}$ , the simulation estimates the density and color of physical entities at position  $\mathbf{x}$  through physically based rendering:  $\sigma_\phi = O_\phi(\mathbf{x}; F_\theta)$ ,  $\mathbf{c}_\phi = \int_{\Omega} L(\mathbf{x}, \omega_i) B_\phi(\mathbf{x}, \mathbf{d}, \omega_i; F_\theta)(\omega_i \cdot N_\phi(\mathbf{x}; F_\theta)) d\omega_i$ , where entity color  $\mathbf{c}_\phi$  is approximated with physically-based rendering equation [29] with normal  $N_\phi$  and BRDF  $B_\phi$ . Importantly, we approximate the incident illumination  $L(\mathbf{x}, \omega_i)$  with radiance by tracing a ray  $\mathbf{r}(t) = \mathbf{x} - t\omega_i$  opposite to the incident direction in the learned NeRF, *i.e.*  $L(\mathbf{x}, \omega_i) = C(\mathbf{r})$ . Depending on the physical entities, we use analytical or sampling-based solutions for the integral. Multiple bounces can be simulated through sampling.



(a) Original NeRF (b) Surface normal (c) Metaball centers (red) (d) Snow with diffuse model (e) Snow with scattering

Figure 6: **Snow simulation.** We first locate metaballs on object surfaces facing upward based on surface normal values (b). With metaballs, we can estimate the density  $\sigma_{snow}(\mathbf{x})$  and color  $\mathbf{c}(\mathbf{x})$  with a parzen window density estimator. (d) and (e) show the differences between fully diffuse model and scattering approximations, shadowed parts in (d) are lit in (e).



Original ClimateGAN [60] Ours

Figure 7: **Smog simulation comparison.** ClimateNeRF simulate smog in view-consistent manner, and separate foreground objects from background better.

We follow the volumetric rendering process defined in two passes. For every point along a camera ray, we query the opacity and color of the physical entities as described in Sec. 3.3. The system also retrieves the original density  $\sigma_\theta$  and color  $\mathbf{c}_\theta$  through Eq. 1. ClimateNeRF estimates final density and color of the simulated scene by:  $\sigma_{\text{final}} = \sigma_\theta + \sigma_\phi$ ,  $\mathbf{c}_{\text{final}} = \frac{\sigma_\theta \mathbf{c}_\theta + \sigma_\phi \mathbf{c}_\phi}{\sigma_\theta + \sigma_\phi}$ , following [35]. Once  $\sigma_{\text{final}}$ ,  $\mathbf{c}_{\text{final}}$  are estimated for each ray points  $\{\mathbf{x}_i\}$ , the pixel color is calculated by volume rendering. Fig. 3 depicts the entire physically inspired simulation and rendering process.

### 3.4. Climate Effect Details

ClimateNeRF is able to simulate smog, flood, and snow through various choices of  $\{O_\phi, N_\phi, B_\phi\}$ .

**Smog Simulation** We assume that smog is formed by tiny absorbing particles, uniformly distributed in empty space. In empty space, the NeRF density  $\sigma_\theta = 0$ . The Beer-Lambert law (originally [5, 36]; in [21]) means we can model smog density in free space by simply adding a non-negative constant to the density. Inside high-density regions of the NeRF, adding the constant does not significantly change the integral, so we have  $O_\phi(\mathbf{x}; F_\theta) = \sigma_{\text{smog}}$ . where  $\sigma_\phi$  decides the density of the smog. Smog particles have a constant color  $\mathbf{c}_{\text{smog}}$ . Both  $\mathbf{c}_{\text{smog}}$  and  $\sigma_{\text{smog}}$  are controllable parameters. Fig. 12 depicts the effects of various smog densities.

**Flood Simulation** The water surface of the flooded scene is approximately a horizontal plane:  $\mathbf{n}_w(\mathbf{x} - \mathbf{o}_w) = 0$ , where the gravity direction normal  $\mathbf{n}_w$  is estimated with camera poses and vanishing points detection [46], and plane origin



(a) Clear water (b) Muddy water

Figure 9: **Refraction in flood simulation**

$\mathbf{o}_w = (0, 0, h)$  determines the water height. But there are water ripples, which we implement following [27] with Fast Fourier Transform (FFT) based ripples and waves. The FFT wave takes random spectral coefficients as input and outputs a spatiotemporal surface normal based on wind speed, direction, and spatial and temporal frequencies. As shown in Fig. 8, compared against still water, FFT-based water surface simulation significantly improves the realism of the water surfaces. We simulate opacity and micro-facet ripples that make the water look glossy (details in supplementary). To approximate the integral of the rendering equation, we adopt the sigma-point method [49, 69] and sample 5 rays from  $\mathbf{x}$ , including reflection direction  $\mathbf{d}_r$  and nearby four rays.

**Snow Simulation** Snow is more likely to be accumulated on surfaces facing upward, and the deeper part of the snow is denser due to gravity. We simulate density over object surfaces using metaballs [52, 32] centered on surfaces and with density  $\sigma_o$  at the center. Realism is ensured by considering both the scene geometry and snowfall direction (e.g., areas such as under the table receiving less accumulated snow). The density distribution within a metaball can be formulated by kernel function  $\mathbf{K}(r, R, \sigma_o)$ , which leads to a smooth decrease of density as the distance  $r$  from the metaball’s center grows. We follow [67] and tuned it for better quality:  $\mathbf{K}(r, R, \sigma) = \frac{315}{64\pi 1.5^7} (1.5^2 - (\frac{r}{R})^2)^3 \sigma$ .

For any point  $\mathbf{x}$  in the space, we calculate the snow’s density of  $\mathbf{x}$  using a parzen window density estimator over  $N$  local nearest neighbors. The final density of the snow surface is decided accordingly to local geometry (details in supplementary). We use a spatially-varying diffuse color  $\mathbf{c}_\phi(\mathbf{x}_i^{(p)})$  (which is close to pure white multiplied by the average illumination of the scene) to approximate BRDF, and apply a subsurface scattering effect [52] to light the snow’s shadowed part (further detail in supplementary). We re-cast shadows on snow by leveraging 2D shadow prediction [12]



Figure 8: **Flood simulation comparison.** Both ClimateGAN++ [60] and Stable Diffusion [55] inpaint flood using water masks from our reconstructions, but suffer from inaccurate reflections or randomly generated contents. ClimateNeRF simulates photorealistic reflections and ripples.

and bake shadow map using NeRF geometry. Fig. 6 depicts the snow simulation process.

## 4. Experiments

We evaluate ClimateNeRF and show simulated results over various scenes across different climate effects. We compare our results with state-of-the-art 2D synthesis and 3D stylization to show the quality and consistency of rendered frames. Experimental results demonstrate that our method is more realistic and faithful than existing 2D synthesis and NeRF model finetuned on stylized images while we also maintain temporal consistency and physical plausibility. We encourage readers to watch supplementary videos for a better demonstration of our method’s quality.

### 4.1. Experimental Details

**Datasets.** We conduct experiments on various outdoor scenes: *Playground*, *Family*, *Horse*, *Truck*, and *Train* from *Tanks and Temples* dataset [34], *Garden* from *Mip-NeRF360* [4] and *Seq 00* from *KITTI-360* [40]. The scenes vary significantly in contents, layouts, and viewpoints.

**Baselines.** We compare ClimateNeRF with the state-of-the-art 2D image editing methods, such as stable diffusion inpainting [55], ClimateGAN [60], as well as state-of-the-art 3D NeRF stylization [78]. For all 2D synthesis approaches, we first build a NeRF using NGP, render at the target view, and conduct synthesis. For all the 3D methods, we re-use the improved version of NGP-based NeRF. **ClimateGAN** [60] uses monocular depth to predict masks and uses GAN to inpaint the climate-related effects, including smog and flood; **ClimateGAN++** is an improved version for flood simulation using our method’s water mask, yielding better geometry consistency; **Swapping Autoencoder** [53] is a photorealistic 2D style transfer method. We use the model pre-trained on Flickr Mountains dataset and Flickr Waterfall dataset [53] for snow. **Stable Diffusion** [55] is the state-of-the-art guided image inpainting method based on latent diffusion model. We feed accurate water masks produced by ClimateNeRF and use text prompts of “flooding” for inpainting. **3D Stylization** leverages FastPhotoStyle [38]. To simulate white snow cov-

erings, we stylize regions labeled *road*, *terrain*, *vegetation*, and *sky* while keeping the geometry. Please refer to supplementary materials for additional implementation details for all competing methods including ours.

### 4.2. Experimental Results

**Qualitative Results** Fig. 7 depicts qualitative results from smog simulation. Our method delivers better realism and physical plausibility (see the different transmission levels across foreground and background). ClimateGAN [60] generates visually reasonable results but fails to provide sharp boundaries. Additionally, video results further show our method is better at view consistency.

We also report flood simulation results in Fig. 8. ClimateGAN++ [60] produces waters with wrong reflection and blurry artifacts. Stable Diffusion [55] provides realistic and diverse colors and reflectance but hallucinates additional contents (e.g., cars, trees) and lacks consistency. ClimateNeRF renders realistic reflections with Fresnel effects and water ripples. Videos show that ClimateNeRF is view-consistent and provides fluid dynamics. Fig. 9 shows that our method simulates both clear and muddy water.

We report snow simulation results in Fig. 10. As the figure shows, 3D Stylization changes the floor texture but cannot add physical entities to the scene, limiting its realism. Swapping Autoencoder [53] changes the overall appearance but hallucinates unrealistic textures (e.g., car texture). On the other hand, ClimateNeRF simulates photorealistic winter effects, including accumulated snow and change of sky and tree colors, etc. ClimateNeRF even piles snow on tiny structures like pedals, as shown in the figure.

Using our proposed NeRF significantly enhances the quality of climate simulation. As demonstrated in Fig. 4, comparing to Instant-NGP, our improved geometry results in a more uniform smog density that is aligned with depth, and a more natural intersection between the border of the water surface and the scene. In addition, the appearance of white floaters is reduced and the ground is more evenly covered by snow in the snow scene.

**User Study** We perform a user study to validate our approach quantitatively. Users are asked to watch pairs of

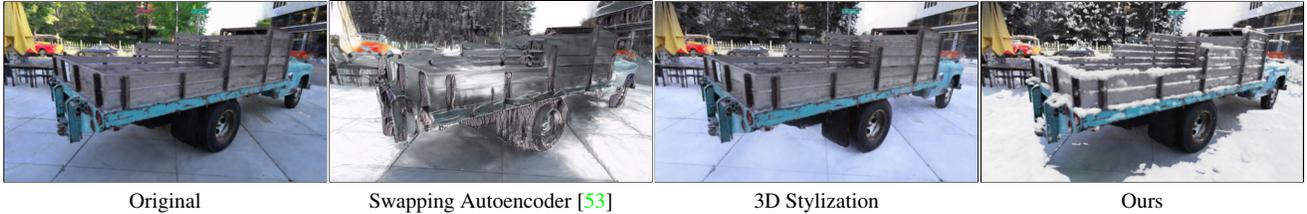


Figure 10: **Snow simulation comparison.** Swapping-Autoencoder [53] captures appearance changes but ignores the geometry of both truck and train. 3D Stylization preserves the geometry of original scene well but doesn't accumulate snow on horizontal surfaces. In contrast, ClimateNeRF has convincing snow accumulation both on ground and on objects. Note small snow accumulations on the bogies and running board on the train and the boards and bonnet of the truck.



Figure 11: **User Study.** The bar length indicates the % of users voting for higher realism. The video quality of ClimateNeRF significantly outperforms all baselines.



Figure 12: **Controllable Simulation.** ClimateNeRF is highly controllable by users. **smog density** (1st row); **flood levels** (2nd row); **snow accumulations** (3rd row).

synthesized images or videos of the same scene and pick the one with higher realism. 37 users participated in the study, and in total, we collected 2664 pairs of comparisons. Results are reported in Fig. 11. ClimateNeRF has consistently been favored among all video simulation comparisons thanks to its high realism and view consistency. Single image comparison does not consider view consistency. In this case, ClimateNeRF still outperforms most baselines except diffusion models on flooding, which also produces realistic water reflectances. Users find diffusion models tend to produce more reflective water surfaces and diverse ripples.



Figure 13: **Simulation on Driving Scenes [40].**

**Rising Sea-level Simulation** ClimateNeRF demonstrates its capability in synthesizing sea-level rise through the flood simulation pipeline outlined in Section 3.4. We selected world-famous scenes from OMMO dataset [45] and Adobe Stock, and simulate the sea-level according to specified height. The water color is set manually to match the original scene. By portraying the potential impacts of global warming on these iconic locations, our intention is to significantly enhance public awareness regarding climate change and advocate for prompt and effective measures.

**Controllability.** One unique advantage of ClimateNeRF is its controllability. Fig. 12 depicts the changes in smog density, water height, and snow thickness. We show additional controllability results, such as ripple size, flood color, water reflectance, and smog color in supp materials.

**Adverse weather simulation for self-driving.** ClimateNeRF can be applied to any NeRF scene. In Fig. 13, ClimateNeRF simulates climate effects in driving scenes of KITTI-360 [40], demonstrating the potential of ClimateNeRF to train and test self-driving under adverse weather.

**Limitations** The simulation performance is closely tied to the quality of NeRF. Despite our proposed improvement, challenging cases might lead to simulation artifacts. Additionally, ClimateNeRF is not real-time currently, and thus acceleration is an area for future work.



Figure 14: **Rising Sea-level Simulation.** We simulate rising sea level at the Sydney Opera House from various view points.



Figure 15: **Rising Sea-level Simulation.** Visualizing the consequences of rising sea levels in world-famous scenes [45] can raise general public awareness about climate change.

## 5. Conclusion

We propose a novel NeRF editing framework that applies physical simulation to NeRF models of scenes. Leveraging this framework, we build ClimateNeRF, allowing us to render realistic climate change effects, including smog, flood, and snow. Our synthesized videos are realistic, view-consistent, physically plausible, and highly controllable. We demonstrate the potential of ClimateNeRF to help raise climate change awareness in the community and enhance self-driving robustness to adverse weather conditions.

**Acknowledgement** This project is partially funded by the Illinois Smart Transportation Initiative STII-21-07, an Amazon Research Award, an Intel Research Gift, an IBM IIDA grant, National Science Foundation under Grant No.2106825 and gifts from Amazon and Boeing. We also thank NVIDIA for the Academic Hardware Grant, and thank Wei-Chiu Ma, Anand Bhattad, and Derek Hoiem for helpful discussion.

## References

- [1] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2stylegan++: How to edit the embedded images? In *CVPR*, 2020. 3
- [2] Benjamin Attal, Jia-Bin Huang, Michael Zollhöfer, Johannes Kopf, and Changil Kim. Learning neural light fields with ray-space embedding. In *CVPR*, pages 19819–19829, 2022. 3
- [3] 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 *ICCV*, 2021. 3, 4
- [4] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *CVPR*, 2022. 2, 3, 4, 5, 7
- [5] Beer. Bestimmung der absorption des rothen lichts in farbigen flüssigkeiten. *Annalen der Physik und Chemie*, 1852. 6
- [6] Sagie Benaim, Frederik Warburg, Peter Ebert Christensen, and Serge Belongie. Volumetric disentanglement for 3d scene manipulation. in *arXiv*, 2022. 3
- [7] Mark Boss, Raphael Braun, Varun Jampani, Jonathan T. Barron, Ce Liu, and Hendrik P.A. Lensch. Nerf: Neural reflectance decomposition from image collections. In *ICCV*, 2021. 3, 4
- [8] Mark Boss, Varun Jampani, Raphael Braun, Ce Liu, Jonathan Barron, and Hendrik Lensch. Neural-pil: Neural pre-integrated lighting for reflectance decomposition. *NeurIPS*, 2021. 3
- [9] Daniel A. Chapman, Adam Corner, Robin Webster, and Ezra M Markowitz. Climate visuals: A mixed methods investigation of public perceptions of climate images in three countries. *Global Environmental Change*, 2016. 1
- [10] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. *ECCV*, 2022. 2, 4

- [11] Yaosen Chen, Qi Yuan, Zhiqiang Li, Yuegen Liu, Wei Wang, Chaoping Xie, Xuming Wen, and Qien Yu. Upst-nerf: Universal photorealistic style transfer of neural radiance fields for 3d scene. *in arXiv*, 2022. 3
- [12] Zhihao Chen, Lei Zhu, Liang Wan, Song Wang, Wei Feng, and Pheng-Ann Heng. A multi-task mean teacher for semi-supervised shadow detection. In *CVPR*, pages 5611–5620, 2020. 6
- [13] Pei-Ze Chiang, Meng-Shiun Tsai, Hung-Yu Tseng, Wei-Sheng Lai, and Wei-Chen Chiu. Stylizing 3d scene via implicit representation and hypernetwork. In *WACV*, 2022. 3
- [14] Chong Bao and Bangbang Yang, Zeng Junyi, Bao Hujun, Zhang Yinda, Cui Zhaopeng, and Zhang Guofeng. Neumesh: Learning disentangled neural mesh-based implicit field for geometry and texture editing. In *ECCV*, 2022. 3
- [15] Blender Online Community. *Blender - a 3D modelling and rendering package*. Blender Foundation, Stichting Blender Foundation, Amsterdam, 2018. 2
- [16] Gautier Cosne, Adrien Juraver, Mélisande Teng, Victor Schmidt, Vahe Vardanyan, Alexandra Luccioni, and Yoshua Bengio. Using simulated data to generate images of climate change. *ICLR Workshop*, 2020. 3
- [17] Epic Games. Unreal engine. 2
- [18] Ronald Fedkiw, Jos Stam, and Henrik Wann Jensen. Visual simulation of smoke. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, 2001. 3
- [19] Bryan E Feldman and James F O’Brien. Modeling the accumulation of wind-driven snow. In *ACM SIGGRAPH conference abstracts and applications*, 2002. 2, 3
- [20] Alain Fournier and William T Reeves. A simple model of ocean waves. In *Proceedings of the 13th annual conference on Computer graphics and interactive techniques*, 1986. 3
- [21] M. Fox. *Optical Properties of Solids (2 ed.)*. Oxford University Press, 2010. 6
- [22] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *CVPR*, 2022. 3
- [23] Erik Glaas, Anne Gammelgaard Ballantyne, Tina-Simone Neset, and Björn-Ola Linnér. Visualization for supporting individual climate change adaptation planning: Assessment of a web-based tool. *Landscape and Urban Planning*, 2017. 1
- [24] John K Haas. A history of the unity game engine. 2014. 2
- [25] Martin Hahner, Dengxin Dai, Christos Sakaridis, Jan-Nico Zaeck, and Luc Van Gool. Semantic understanding of foggy scenes with purely synthetic data. In *IEEE Intelligent Transportation Systems Conference (ITSC)*, 2019. 3
- [26] Jamie Herring, Matthew S VanDyke, R Glenn Cummins, and Forrest Melton. Communicating local climate risks online through an interactive data visualization. *Environmental Communication*, 2017. 1
- [27] Christopher J Horvath. Empirical directional wave spectra for computer graphics. In *Proceedings of the 2015 Symposium on Digital Production*, 2015. 2, 3, 6
- [28] Liming Jiang, Changxu Zhang, Mingyang Huang, Chunxiao Liu, Jianping Shi, and Chen Change Loy. Tsit: A simple and versatile framework for image-to-image translation. In *ECCV*, 2020. 3
- [29] James T Kajiya. The rendering equation. In *Proceedings of the 13th annual conference on Computer graphics and interactive techniques*, 1986. 5
- [30] James T Kajiya and Brian P Von Herzen. Ray tracing volume densities. *ACM SIGGRAPH computer graphics*, 1984. 4
- [31] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019. 3
- [32] Leonid Keselman and Martial Hebert. Approximate differentiable rendering with algebraic surfaces. In *ECCV*, 2022. 6
- [33] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM TOG*, 2017. 2
- [34] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM TOG*, 2017. 7
- [35] Sosuke Kobayashi, Eiichi Matsumoto, and Vincent Sitzmann. Decomposing nerf for editing via feature field distillation. In *NeurIPS*, 2022. 3, 6
- [36] J.H. Lambert. *Photometria sive de mensura et gradibus luminis, colorum et umbrae*. Eberhardt Klett, 1760. 6
- [37] Zoe Leviston, Jennifer Price, and Brian Bishop. Imagining climate change: The role of implicit associations and affective psychological distancing in climate change responses. *European Journal of Social Psychology*, 2014. 1
- [38] Yijun Li, Ming-Yu Liu, Xueting Li, Ming-Hsuan Yang, and Jan Kautz. A closed-form solution to photorealistic image stylization. In *ECCV*, 2018. 5, 7
- [39] Zhengqin Li, Jia Shi, Sai Bi, Rui Zhu, Kalyan Sunkavalli, Miloš Hašan, Zexiang Xu, Ravi Ramamoorthi, and Manmohan Chandraker. Physically-based editing of indoor scene lighting from a single image. *ECCV*, 2022. 3
- [40] Yiyi Liao, Jun Xie, and Andreas Geiger. Kitti-360: A novel dataset and benchmarks for urban scene understanding in 2d and 3d. *IEEE TPAMI*, 2022. 2, 7, 8
- [41] Zhi-Hao Lin, Wei-Chiu Ma, Hao-Yu Hsu, Yu-Chiang Frank Wang, and Shenlong Wang. Neurmips: Neural mixture of planar experts for view synthesis. In *CVPR*, 2022. 2, 3
- [42] Huan Ling, Karsten Kreis, Daiqing Li, Seung Wook Kim, Antonio Torralba, and Sanja Fidler. Editgan: High-precision semantic image editing. In *NeurIPS*, 2021. 3
- [43] Difan Liu, Sandesh Shetty, Tobias Hinz, Matthew Fisher, Richard Zhang, Taesung Park, and Evangelos Kalogerakis. Asset: autoregressive semantic scene editing with transformers at high resolutions. *ACM TOG*, 2022. 3
- [44] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. *NeurIPS*, 2020. 3
- [45] Chongshan Lu, Fukun Yin, Xin Chen, Tao Chen, Gang Yu, and Jiayuan Fan. A large-scale outdoor multi-modal dataset and benchmark for novel view synthesis and implicit scene reconstruction. *in arXiv*, 2023. 8, 9
- [46] Xiaohu Lu, Jian Yaoy, Haoang Li, Yahui Liu, and Xiaofeng Zhang. 2-line exhaustive searching for real-time vanishing

- point estimation in manhattan world. In *WACV*. IEEE, 2017. 6
- [47] Julien NP Martel, David B Lindell, Connor Z Lin, Eric R Chan, Marco Monteiro, and Gordon Wetzstein. Acorn: Adaptive coordinate networks for neural scene representation. *ACM TOG*, 2021. 3
- [48] Ricardo Martin-Brualla, Noha Radwan, Mehdi SM Sajjadi, Jonathan T Barron, Alexey Dosovitskiy, and Daniel Duckworth. Nerf in the wild: Neural radiance fields for unconstrained photo collections. In *CVPR*, 2021. 5
- [49] Henrique MT Menegaz, João Y Ishihara, Geovany A Borges, and Alessandro N Vargas. A systematization of the unscented kalman filter theory. *IEEE Transactions on automatic control*, 2015. 6
- [50] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. 1, 2, 3, 4, 5
- [51] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM TOG*, 2022. 1, 2, 3, 4, 5
- [52] Tomoyuki Nishita, Hiroshi Iwasaki, Yoshinori Dobashi, and Eihachiro Nakamae. A modeling and rendering method for snow by using metaballs. In *Computer Graphics Forum*, 1997. 3, 6
- [53] Taesung Park, Jun-Yan Zhu, Oliver Wang, Jingwan Lu, Eli Shechtman, Alexei Efros, and Richard Zhang. Swapping autoencoder for deep image manipulation. *NeurIPS*, 2020. 3, 7, 8
- [54] Christian Reiser, Songyou Peng, Yiyi Liao, and Andreas Geiger. Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. In *ICCV*, 2021. 3
- [55] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 2, 3, 7
- [56] Viktor Rudnev, Mohamed Elgharib, William Smith, Lingjie Liu, Vladislav Golyanik, and Christian Theobalt. Neural radiance fields for outdoor scene relighting. *ECCV*, 2022. 3
- [57] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. *in arXiv*, 2022. 3
- [58] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, et al. Photorealistic text-to-image diffusion models with deep language understanding. *in arXiv*, 2022. 3
- [59] Victor Schmidt, Alexandra Luccioni, S Karthik Mukkavilli, Narmada Balasooriya, Kris Sankaran, Jennifer Chayes, and Yoshua Bengio. Visualizing the consequences of climate change using cycle-consistent adversarial networks. *ICLR*, 2019. 3
- [60] Victor Schmidt, Alexandra Sasha Luccioni, Mélisande Teng, Tianyu Zhang, Alexia Reynaud, Sunand Raghupathi, Gautier Cosne, Adrien Juraver, Vahe Vardanyan, Alex Hernandez-Garcia, et al. Climategan: Raising climate change awareness by generating images of floods. *ICLR*, 2022. 2, 3, 6, 7
- [61] Pratul P Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, and Jonathan T Barron. Nerv: Neural reflectance and visibility fields for relighting and view synthesis. In *CVPR*, 2021. 3, 4
- [62] Alexey Stomakhin, Craig Schroeder, Lawrence Chai, Joseph Teran, and Andrew Selle. A material point method for snow simulation. *ACM TOG*, 2013. 2, 3
- [63] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *CVPR*, 2022. 3, 4
- [64] Matthew Tancik, Vincent Casser, Xinchen Yan, Sabeek Pradhan, Ben Mildenhall, Pratul P Srinivasan, Jonathan T Barron, and Henrik Kretzschmar. Block-nerf: Scalable large scene neural view synthesis. In *CVPR*, 2022. 2
- [65] Matthew Tancik, Ethan Weber, Evonne Ng, Ruilong Li, Brent Yi, Justin Kerr, Terrance Wang, Alexander Kristoffersen, Jake Austin, Kamyar Salahi, et al. Nerfstudio: A modular framework for neural radiance field development. *in arXiv*, 2023. 5
- [66] Jerry Tessendorf. Simulating ocean water. In *SIGGRAPH*, 2001. 3
- [67] Kees van Kooten, Gino van den Bergen, and Alex Telea. *Point-based visualization of metaballs on a gpu*. Addison-Wesley Longman, 2007. 6
- [68] Dor Verbin, Peter Hedman, Ben Mildenhall, Todd Zickler, Jonathan T. Barron, and Pratul P. Srinivasan. Ref-NeRF: Structured view-dependent appearance for neural radiance fields. *CVPR*, 2022. 3, 4
- [69] Eric A Wan and Rudolph Van Der Merwe. The unscented kalman filter for nonlinear estimation. In *Proceedings of the IEEE Adaptive Systems for Signal Processing, Communications, and Control Symposium*, 2000. 6
- [70] Can Wang, Menglei Chai, Mingming He, Dongdong Chen, and Jing Liao. Clip-nerf: Text-and-image driven manipulation of neural radiance fields. In *CVPR*, 2022. 3
- [71] Suttisak Wizadwongsa, Pakkapon Phongthawee, Jiraphon Yenphraphai, and Supasorn Suwajanakorn. Nex: Real-time view synthesis with neural basis expansion. In *CVPR*, 2021. 3
- [72] Liwen Wu, Jae Yong Lee, Anand Bhattad, Yu-Xiong Wang, and David Forsyth. Diver: Real-time and accurate neural radiance fields with deterministic integration for volume rendering. In *CVPR*, 2022. 2, 3
- [73] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. In *NeurIPS*, 2021. 4
- [74] Qiangeng Xu, Zexiang Xu, Julien Philip, Sai Bi, Zhixin Shu, Kalyan Sunkavalli, and Ulrich Neumann. Point-nerf: Point-based neural radiance fields. In *CVPR*, 2022. 3
- [75] Tianhan Xu and Tatsuya Harada. Deforming radiance fields with cages. In *ECCV*, 2022. 3
- [76] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. Plenotrees for real-time rendering of neural radiance fields. In *ICCV*, 2021. 3

- [77] Yu-Jie Yuan, Yang-Tian Sun, Yu-Kun Lai, Yuewen Ma, Rongfei Jia, and Lin Gao. Nerf-editing: Geometry editing of neural radiance fields. In *CVPR*, 2022. 3
- [78] Kai Zhang, Nick Kolkin, Sai Bi, Fujun Luan, Zexiang Xu, Eli Shechtman, and Noah Snavely. Arf: Artistic radiance fields. *ECCV*, 2022. 2, 3, 7
- [79] Kai Zhang, Fujun Luan, Qianqian Wang, Kavita Bala, and Noah Snavely. Physg: Inverse rendering with spherical gaussians for physics-based material editing and relighting. In *CVPR*, 2021. 3
- [80] Qiang Zhang, Seung-Hwan Baek, Szymon Rusinkiewicz, and Felix Heide. Differentiable point-based radiance fields for efficient view synthesis. *in arXiv*, 2022. 3
- [81] Xiuming Zhang, Pratul P Srinivasan, Boyang Deng, Paul Debevec, William T Freeman, and Jonathan T Barron. Nerfactor: Neural factorization of shape and reflectance under an unknown illumination. *ACM TOG*, 2021. 3
- [82] Shuaifeng Zhi, Tristan Laidlow, Stefan Leutenegger, and Andrew Davison. In-place scene labelling and understanding with implicit scene representation. In *ICCV*, 2021. 5
- [83] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017. 3
- [84] Peiye Zhuang, Oluwasanmi Koyejo, and Alexander G Schwing. Enjoy your editing: Controllable gans for image editing via latent space navigation. In *ICLR*, 2021. 3
- [85] Károly Zsolnai-Fehér. The flow from simulation to reality. *Nature Physics*, 2022. 2
- [86] Yiming Zuo and Jia Deng. View synthesis with sculpted neural points. *in arXiv*, 2022. 3