Nexels: Neurally-Textured Surfels
for Real-Time Novel View Synthesis with Sparse Geometries

Victor Rong[*1,2]   Jan Held[*3,4]   Victor Chu[1,2]   Daniel Rebain[5]   Marc Van Droogenbroeck[4]
Kiriakos N. Kutulakos[1,2]   Andrea Tagliasacchi[†1,3]   David B. Lindell[†1,2]

[1]University of Toronto [2]Vector Institute [3]Simon Frasier University [4]University of Liège [5]University of British Columbia

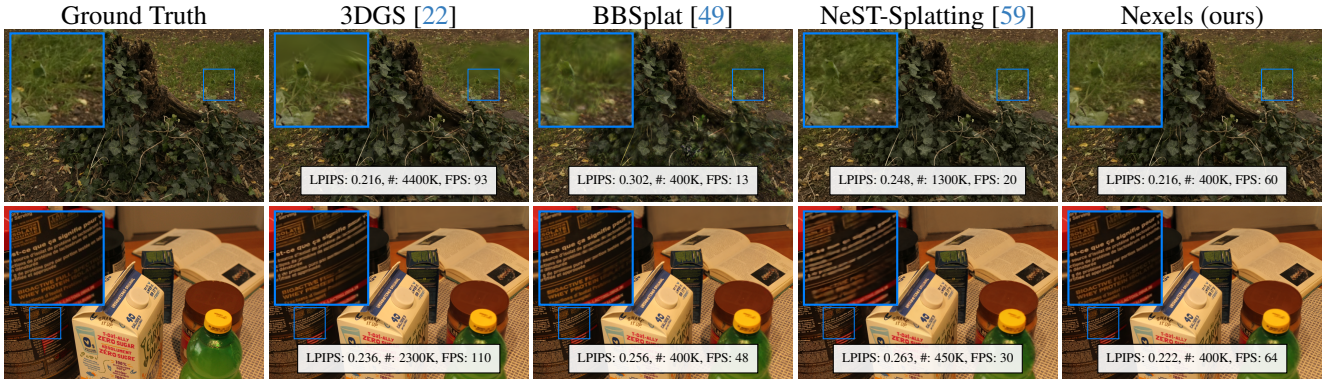[*,†]equal contribution

https://lessvrong.com/cs/nexels

Figure 1. **Teaser.** Our nexels match the quality of 3D Gaussian splats with a fraction of the number of primitives, and at faster rendering speeds than prior textured primitive works. We show test-set averages of the LPIPS↓, # of primitives, and frames-per-second for two scenes: the STUMP scene from the Mip-NeRF360 dataset [2] and the GROCERY scene from our dataset.

Abstract

Though Gaussian splatting has achieved impressive results in novel view synthesis, it requires millions of primitives to model highly textured scenes, even when the geometry of the scene is simple. We propose a representation that goes beyond point-based rendering and decouples geometry and appearance in order to achieve a compact representation. We use surfels for geometry and a combination of a global neural field and per-primitive colours for appearance. The neural field textures a fixed number of primitives for each pixel, ensuring that the added compute is low. Our representation matches the perceptual quality of 3D Gaussian splatting while using $9.7\times$ fewer primitives and $5.5\times$ less memory on outdoor scenes and using $31\times$ fewer primitives and $3.7\times$ less memory on indoor scenes. Our representation also renders twice as fast as existing textured primitives while improving upon their visual quality.

1. Introduction

In computer graphics, a scene's geometry governs how light is blocked, while its appearance determines the color contributed to each rendered pixel. Point-based representations like 3D Gaussian splats (3DGS) *merge* these roles, since each primitive jointly encodes geometry and appearance [22]. This coupling stands in contrast to traditional surface-based methods, where coarse geometry (e.g., a quad) can still capture fine-grained appearance (e.g., a detailed painting) through texturing. As illustrated in Figure 2, Gaussian splats therefore require a large number of primitives to represent high-frequency appearance, leading to substantial memory usage.

A more compact representation becomes possible by *decoupling* geometry from appearance. When appearance is delegated to textures, mesh-based models can depict highly realistic scenes, as commonly seen in professionally produced games and films. Yet meshes have lagged behind in the task of novel view synthesis, where the objective is to generate new viewpoints from a collection of input images of a 3D scene. In the dense-pose setting we consider, the dominant strategy is *differentiable rendering*, which optimizes scene parameters (geometry, light, and materials) so that rendered views match the training images. Within this framework, meshes face difficulties due to their discrete connectivity, potential for self-intersections, and complica-
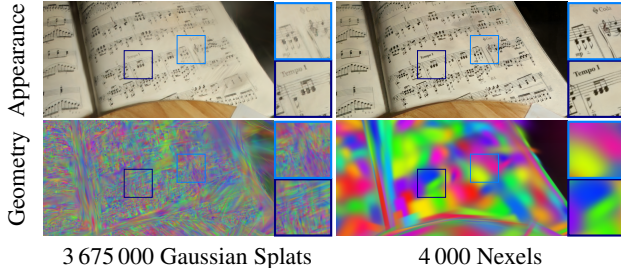
1

Figure 2. **Limitation of point-based representations.** Our captured TABLE scene contains planar geometries with complex textures such as the sheet music above. An ideal representation would have very few parameters devoted to geometry, and much more devoted to appearance. Gaussian splatting needs millions of Gaussians to represent the fine details, each of which has a *fixed* ratio of geometry-to-appearance parameters.

tions in computing silhouette gradients [26, 29, 31, 37].

To mitigate the need for an excessive number of primitives without resorting to the complexities of mesh optimization, recent work has proposed decoupling geometry and appearance in Gaussian splatting by introducing *textured splats*, most commonly via per-primitive image textures [7, 38, 42, 49]. However, the memory footprint of such textures grows quadratically with resolution, which constrains their ability to represent high-frequency appearance in large-scale scenes. Zhang et al. [59] attempt to overcome this limitation by employing a neural field to provide appearance in a deferred manner: each primitive samples neural features from a multi-resolution hash grid, which are then alpha-composited and passed through a feed-forward network. While this design requires only a single network evaluation, hash-grid lookups are similarly expensive, and performing several of them for every ray–primitive intersection results in *slow rendering* (<30 FPS), as well as *slow training*.

Although these methods take inspiration from textured meshes, they miss a crucial observation: in standard mesh rasterization with depth buffering, texture fetches are performed only for a *limited* set of *opaque* surface fragments [1, 21]. Motivated by this, we introduce a new representation that preserves the optimization advantages of point-based approaches while leveraging the lightweight appearance modeling of mesh-based methods – all without sacrificing real-time rendering performance.

Contributions. We introduce *nexels*, a novel neurally-textured primitive that overcomes the coupling of geometry and appearance in Gaussian splatting. In particular: (i) For geometry, we introduce a differentiable *quad* indicator, enabling us to better reconstruct surfaces and sharp boundaries; (ii) For appearance, we learn a world-space neural field that provides view-dependent texture only for the most relevant primitives, capturing fine details efficiently while keeping computation low; (iii) We introduce a new dataset that highlights the limitations of point-based representations in accurately reconstructing regions with high--frequency textures; (iv) We achieve perceptual parity with 3DGS using $9.7\times$ fewer primitives on outdoor scenes and $31\times$ fewer primitives on indoor scenes; (v) Compared to concurrent texture methods, we achieve better photometric quality, while rendering more than twice as fast.

2. Related Work

2.1. Neural Fields

A neural field implicitly represents a quantity over a region, such as 3D space, through neural network queries. In novel view synthesis, Mildenhall et al. [34] use neural fields to parameterize per-point radiance and density for volumetric rendering. These neural radiance fields (NeRFs) have slow rendering speed due to both the large number of samples needed for volumetric rendering as well as expensive per-sample calculations. To reduce the number of samples used for rendering volumes, early NeRF papers employ stratified sampling and auxiliary proposal networks to improve the sampling efficiency [2, 3, 34]. Later works use empty space skipping and other acceleration techniques to avoid unnecessary queries [28, 36, 46]. The number of queries can be further constrained using explicit representations such as shells and meshes [45, 52]. In parallel, follow-up methods reduce the computational cost of each query by moving much of the neural network's representational power into spatial features backed by voxel grids [13, 48], hierarchical grids [33, 36, 50, 57], point sets [9, 10, 53], and other structures [6, 8]. Another line of work avoids the slow runtime of neural fields by only using them in an initial stage of training before baking in the appearance and geometry into a fast mesh representation [14, 40, 51, 55]. In contrast, we use surfels to obtain sample locations which avoids the added complication of mesh extraction while seamlessly integrating into the differentiable rendering scheme.

2.2. Gaussian Splatting

Gaussian splatting completely forgoes the neural network and instead represents radiance fields as a set of explicit primitives which have decaying opacities according to a Gaussian distribution [22]. It leverages fast point-based rasterization schemes to achieve real-time performance. There has been a deluge of follow-up works, of which we will only focus on those that present new geometry or appearance representations. Several works find that using flattened 3D primitives also results in high-quality novel view synthesis results [11, 20] and Ye et al. [56] use these to explicitly model opaque surfaces. Other non-Gaussian geometries have also been found to be effective: beta kernels, polyno-

Method	Scene	Ada	Speed	Mem
Texture-GS [54]	✗	✗	✗	✗
GStex [42]	✓	✗	✓	✗
Textured Gaussians [7]	✓	✗	✗	✗
BBSplat [49]	✓	✓	✗	✗
NeST-Splatting [59]	✓	✗	✗	✓
GS-Texturing [38]	✓	✓	✓	✗
Nexels	✓	✓	✓	✓

Table 1. **Prior and concurrent textured works.** We summarize the limitations of recent representations with textured primitives. The criteria are *scene-level reconstructions*, *adaptive density control with textures*, *real-time rendering speeds (30+ FPS)*, and *ability to capture texture details in a memory-efficient manner*.

Diffuse ⟶ Sharp

Figure 3. **Kernel.** We use a generalized Gaussian kernel in order to model near-opaque primitives. Values below 0.1 are set to 0 for visualization purposes. (Left) For $\gamma{=}1$, our nexels correspond to a Gaussian distribution. (Right) As $\gamma{\to}\infty$, they converge toward the indicator of a quad.

mial kernels, triangles, and more have been proposed [15–18, 30, 35]. Beta splats have been particularly effective at achieving higher rendering quality with fewer parameters. They achieve this reduction through a leaner set of parameters for view-dependent appearance, which is orthogonal to our own work [30]. A number of Gaussian splatting works use neural fields in various ways, but unlike our work, the splats in these works have no spatial variation across them and cannot represent details smaller than themselves [32, 41].

A recent line of work applies textures to Gaussian splats to increase the appearance capacity of the representation. This has been done through per-primitive image textures [7, 38, 42, 49] and neural fields [54, 59]. Despite promising results, there are three key issues which hinder these methods, summarized in Tab. 1. First, the majority of these works initialize from a trained dense Gaussian splatting model, which limits their ability to optimize for sparse geometries. Furthermore, the texturing in all prior methods is applied for every ray-primitive interaction. For BBSplat and Textured Gaussians, whose alpha textures further increase overdraw, renders of scene-level reconstructions typically take over fifty milliseconds. Finally, an ideal texture should be able to capture multiscale details within a reasonable amount of memory. Works which use per-primitive image textures are limited by their resolution, which scales poorly with memory. As a partial solution, concurrent work from Papantonakis et al. [38] use image textures with adaptive resolutions. Only NeST-Splatting fully overcomes the issue by using a neural field, but this results in much slower renders. Our representation uses a neural field at a fixed number of samples per-pixel, leading to real-time renders.

3. Preliminaries

Our work builds on both Gaussian splatting and neural field methods. In particular, we make use of the differentiable surfel rasterization of 2D Gaussian splatting (2DGS) [20]

and the neural field architecture of Instant-NGP [36].

2D Gaussian Splatting. Each surfel in 2D Gaussian splatting (2DGS) is a 2D Gaussian in 3D space parameterized by its opacity $o \in \mathbb{R}$, rotation matrix $\mathbf{R} = [\mathbf{v_1}, \mathbf{v_2}, \mathbf{v_3}] \in \mathbb{SO}_3$, mean $\boldsymbol{\mu} \in \mathbb{R}^3$, and scale $\boldsymbol{\sigma} \in \mathbb{R}^2$. The first two axes of the rotation, $\mathbf{v_1}$ and $\mathbf{v_2}$, determine the axes of the Gaussian, while $\mathbf{v_3}$ functions as the surfel normal. Additionally, each primitive has a view-dependent radiance $\mathbf{c}^{\odot} \in \mathbb{R}^{3S}$ parameterized by channel-wise spherical harmonics coefficients, where S is the number of coefficients for a predetermined degree.

Following 3DGS, the surfels are preprocessed and sorted by the depth of their centers. For a ray $\mathbf{r} := \mathbf{o} + t\mathbf{d}$ emanating from the camera position, the intersection in primitive space (u, v) can be computed differentiably from its geometry parameters [20, 47]. This is used to compute its alpha value, $\alpha := o \exp\left(-(u^2 + v^2)/2\right)$. The primitives are alpha-composited in a tile-based rasterizer. The final color is

$$\mathbf{c}(\mathbf{r}) := \sum_i T_i \alpha_i \mathbf{c}_i^{\odot}(\mathbf{r}), \quad T_i := \prod_j^{i-1}(1 - \alpha_i). \quad (1)$$

Instant-NGP. Instant-NGP is a neural field architecture composed of a multiresolution hash-grid $\mathcal{H}$ and a tiny multi-layer perceptron (MLP) network $\mathcal{G}$. The multiresolution hash-grid has L levels, each defining a hash-table $\mathcal{H}_\ell$ at given scale s_ℓ. In our implementation, each hash-table has exactly T entries. Finally, it has a feature dimension of F. Altogether it has $L \times T \times F$ parameters.

Given an input coordinate $x \in \mathbb{R}^3$, an Instant-NGP $\mathcal{F}$ can be configured to return a vector whose size is the MLP output dimension. First, a hash $\mathbf{h}_\ell \in \{0, 1, \dots, T - 1\}$ is calculated for each hash-grid level ℓ based on $s_\ell \cdot x$. This hash is indexed into the grid at that level to retrieve a feature vector. These are concatenated to obtain

$$\mathcal{H}(x) := \left(\mathcal{H}_{0,\mathbf{h}_0}, \dots, \mathcal{H}_{L-1,\mathbf{h}_{L-1}}\right) \in \mathbb{R}^{L \times F}. \quad (2)$$

Finally, the MLP is applied to yield $\mathcal{F}(x) := \mathcal{G}(\mathcal{H}(x))$.
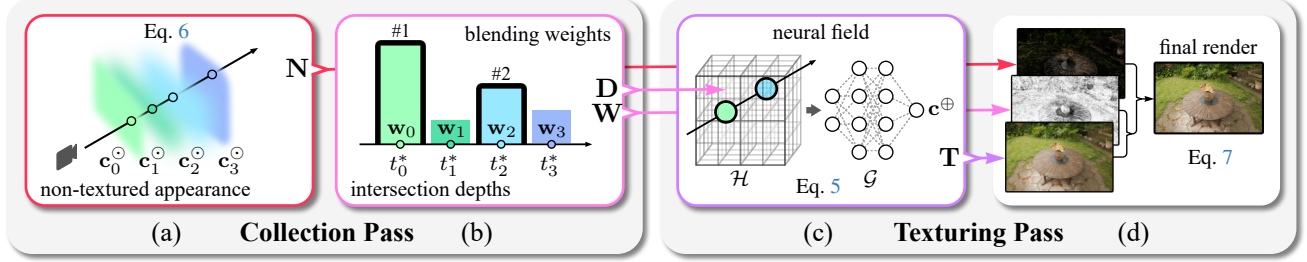
Figure 4. **Rendering pipeline.** (a) We iterate over the surfels and determine the weights at each intersection while also returning a non-textured render. (b) We collect the K intersections which have the largest weights. (c) We pass their K positions into a neural field to obtain textured appearance. (d) We composite the textures with the non-textured render to get the final output.

4. Nexels

This section provides an overview of the main components of our method. Section 4.1 introduces the geometry and appearance representations, while Section 4.2 describes the rendering process. Finally, Section 4.3 outlines the optimization framework, including adaptive density control and the loss terms used during training. For notational convenience, we present our method for a single pixel p with ray $\mathbf{r}$ and use view-*independent* radiance in equations.

4.1. Representation

Geometry. We represent the geometry of nexels using a set of N surfels defined by a kernel governing their opacities. Following 2DGS, the i^{th} primitive has a mean $\mu_i \in \mathbb{R}^3$, rotation $\mathbf{R}_i \in \mathbb{SO}_3$, scale $\sigma_i \in \mathbb{R}^2$, and opacity $\boldsymbol{o}_i \in \mathbb{R}$. In addition, we have a gamma parameter $\gamma_i \in \mathbb{R}^2$ that controls how rect-like the primitive is. The opacity at an intersection (u, v) in primitive space is defined by the generalized Gaussians,

$$\alpha := \boldsymbol{o}_i \exp\left(-\frac{|u|^{2\gamma_{i,1}}}{2}\right) \exp\left(-\frac{|v|^{2\gamma_{i,2}}}{2}\right). \quad (3)$$

Figure 3 shows the nexel shape for different values of γ. As $\gamma \to 1$, the kernel converges to a standard Gaussian. As $\gamma \to \infty$ it converges toward a quad with a sharp transition. This kernel is well suited for modeling opaque and flat surfaces: for large $\boldsymbol{o}$ and γ, the function remains close to 1 over most of the support, unlike a conventional Gaussian which has a long tail. Our kernel behaves like a Gaussian where smoothness is needed, while still being able to produce hard edges in regions with sharp transitions. We constrain $\gamma \geq 1$, which guarantees that our kernel is C_∞. Taking inspiration from the separable filters of signal processing, we choose to take the outer product between the u and v axes. As γ grows, the kernel approaches a rectangle rather than an ellipse as in prior work [15, 30].

Appearance. Each surfel has two modes of appearance: non-textured and textured. During rendering, we select which to use on a per-pixel basis. The non-textured appearance of surfel i is the same as 2DGS [20], $\mathbf{c}_i^\odot \in \mathbb{R}^{3S}$.

The textured appearance of each primitive is represented using a shared neural field $\mathcal{F}$ backed by an Instant-NGP architecture [36]. The field takes as input any 3D position $x \in \mathbb{R}^3$ and outputs a view-dependent radiance $\mathcal{F}(x)$ as a set of spherical harmonics coefficients. The texture at a given point on any of the surfels can be computed by querying the neural field at that position.

For a primitive i, we define $t_i^*(\mathbf{r})$ as the intersection depth and $x_i^*(\mathbf{r}) := \mathbf{o} + t^*\mathbf{d}$ as the world-space intersection. Naively, the textured appearance for pixel p would be $\mathcal{F}(x_i^*(\mathbf{r}))$. Using point samples for differentiable rendering leads to inaccuracies in the fine levels of the hash-grid. Following the down-weighting analysis of Barron et al. [4], we multiply the hash-grid features of level ℓ by

$$\Delta_{i,\ell}(\mathbf{r}) := 1 - \exp\left(-\frac{1}{2\pi}\left(\frac{f}{s_\ell t_i^*(\mathbf{r})}\right)^2\right), \quad (4)$$

where f is the focal length. This requires the same amount of computation as the naive sampling and removes noticeable aliasing artifacts. Ultimately, we define the filtered texture radiance along $\mathbf{r}$ for primitive i as

$$\mathbf{c}_i^\oplus(\mathbf{r}) := \mathcal{G}\left(\Delta_i(\mathbf{r}) \cdot \mathcal{H}(x_i^*(\mathbf{r}))\right). \quad (5)$$

Note that the only primitive information needed to compute $\mathbf{c}_i^\oplus(\mathbf{r})$ is the intersection depth $t_i^*(\mathbf{r})$.

4.2. Rendering

Our rendering step is composed of *two passes* over the image pixels. The first pass computes the initial render using only the non-textured appearance and collects information on which primitives should be textured. The second pass applies the neural texture at the collected ray-primitive intersections and updates the initial render. Figure 4 provides an overview of the different rendering passes.

Collection Pass. In the first pass, we compute the alpha α_i and transmittance values T_i across the M surfels which

4

intersect a pixel's ray. We use eq. (1) to compute the initial render $\mathbf{N}$ for each pixel, assuming no primitive is textured.

When rendering a single pixel, we limit the number of primitives which are textured to a hyperparameter $K \ll M$. We determine the subset of primitives which are textured based on their blending weights which we denote as $\mathbf{w}_i := \alpha_i T_i$. For a given pixel, the K primitives with the highest weights are selected. Inspired by fragment buffer techniques [5], we maintain a per-pixel buffer of size K. Each buffer entry contains a primitive id, its weight, and its depth, totalling to $3K$ 32-bit registers. Initially, the buffer is empty. When iterating through the M primitives and computing the weights, we update the buffer to store the running primitives with the highest weights. The final id, depth, and weight buffers are written into $H \times W \times K$ images $\mathbf{I}, \mathbf{D}$, and $\mathbf{W}$, respectively. These are returned along with the non-textured render $\mathbf{N}$. Additionally, for the K primitives in $\mathbf{I}$, we remove their non-textured colour from the render $\mathbf{N}$. The final value of the non-textured render at pixel p is

$$\mathbf{N}_p := \sum_{i=0}^{M-1} [i \notin \mathbf{I}_p] \, \mathbf{w}_i \mathbf{c}_i^{\odot}. \qquad (6)$$

Texturing Pass. The texturing pass first computes world-space intersection positions from the depth-map buffer $\mathbf{D}$, yielding $x \in \mathbb{R}^{H \times W \times K \times 3}$. Following eq. (5), we query the neural field to compute the filtered texture $\mathbf{T} \in \mathbb{R}^{H \times W \times K \times 3}$ where $\mathbf{T}_{p,j} := \mathbf{c}_{\mathbf{I}_{p,j}}^{\oplus}(\mathbf{r})$ for pixel p with corresponding ray $\mathbf{r}$. Extending eq. (1) to include the choice between texturing and not texturing, we have

$$\sum_{i=0}^{M-1} T_i \alpha_i \mathbf{c}_i(\mathbf{r}) = \sum_{i=0}^{M-1} \mathbf{w}_i \left([i \notin \mathbf{I}_p] \, \mathbf{c}_i^{\odot} + [i \in \mathbf{I}_p] \, \mathbf{c}_i^{\oplus}(\mathbf{r}) \right)$$

$$= \mathbf{N}_p + \sum_{j=0}^{K-1} \mathbf{W}_{p,j} \mathbf{T}_{p,j}. \qquad (7)$$

Hence, we compute the final render from $\mathbf{N}, \mathbf{W}$, and $\mathbf{T}$.

4.3. Optimization

We follow the scheme of 3DGS [22] and optimize against training views for 30,000 iterations with regularizations. We also interleave density control operations.

Adaptive Density Control. We introduce a hyperparameter P to control the number of primitives. Similar to BBSplat [49], we sample the point cloud output of COLMAP for initialization [44]. Specifically, we use farthest point sampling to reduce the initial point cloud to $0.5P$ points [27]. We initialize the scale at each point based on its distance to its predecessors during farthest point sampling.

We also prune unnecessary primitives and add new ones in regions where the current geometry is insufficient. We

perform a densification and pruning step every 100 iterations from iteration 500 to 25,000. We select 5% of primitives to split evenly along their longest axis. The selection is stochastic with probabilities defined using blended errors similar to Rota Bulò et al. [43]'s method. We cap the number of primitives added in the splitting by P. Afterward, we prune all primitives whose opacity is below 0.005.

Losses. At each iteration, we render the nexels from a given training view to produce a predicted image. Following Kerbl et al. [22], we compute a photometric loss between the prediction and ground truth using a mixture of an L_1 loss $\mathcal{L}_1$ and D-SSIM loss $\mathcal{L}_{\text{D-SSIM}}$ [22],

$$\mathcal{L}_{\text{image}} := (1 - \lambda_{\text{D-SSIM}})\mathcal{L}_1 + \lambda_{\text{D-SSIM}}\mathcal{L}_{\text{D-SSIM}}. \qquad (8)$$

We also include a texture loss, $\mathcal{L}_{\text{texture}}$, so that the non-premultiplied texture fits to the ground truth image I^{GT}:

$$\frac{1}{3HW} \sum_{y=0}^{H-1} \sum_{x=0}^{W-1} \sum_{c=0}^{2} \left| I_{y,x,c}^{\text{GT}} - \frac{\sum_{j=0}^{K-1} \mathbf{W}_{y,x,j} \mathbf{T}_{y,x,j,c}}{\sum_{j=0}^{K-1} \mathbf{W}_{y,x,j}} \right|. \qquad (9)$$

To penalize non-textured rendering, we minimize the total blending weights of the non-textured appearance,

$$\mathcal{L}_{\text{alpha}} := \frac{1}{HW} \sum_{y=0}^{H-1} \sum_{x=0}^{W-1} \left(1 - \sum_{j=0}^{K-1} \mathbf{W}_{y,x,j} \right). \qquad (10)$$

We include an L_1 loss on the opacity to induce sparsity [23],

$$\mathcal{L}_{\text{opacity}} := \frac{1}{N} \sum_{i=0}^{N-1} \boldsymbol{o}_i. \qquad (11)$$

Finally, we add a grid weight regularization term,

$$\mathcal{L}_{\text{grid}} := \sum_{\ell=0}^{L-1} \sum_{i=0}^{T-1} \sum_{j=0}^{F-1} s_i^{-3} \mathcal{H}_{\ell,i,j} \qquad (12)$$

which encourages zero-mean grid values that justify the down-weighting done for grid anti-aliasing [4]. The final loss is

$$\mathcal{L} := \mathcal{L}_{\text{image}} + \lambda_{\text{alpha}}\mathcal{L}_{\text{alpha}} + \lambda_{\text{texture}}\mathcal{L}_{\text{texture}} + \lambda_{\text{opacity}}\mathcal{L}_{\text{opacity}} + \lambda_{\text{grid}}\mathcal{L}_{\text{grid}}. \qquad (13)$$

We set $\lambda_{\text{alpha}} = 0.005$, $\lambda_{\text{texture}} = 0.5$, $\lambda_{\text{opacity}} = 0.01$, $\lambda_{\text{grid}} = 0.01$, and $\lambda_{\text{D-SSIM}} = 0.2$ for all experiments.

5. Results

Our method is implemented in Python and CUDA on top of the INRIA 3DGS codebase [22]. We use PyTorch for performing parameter activations and composing the render
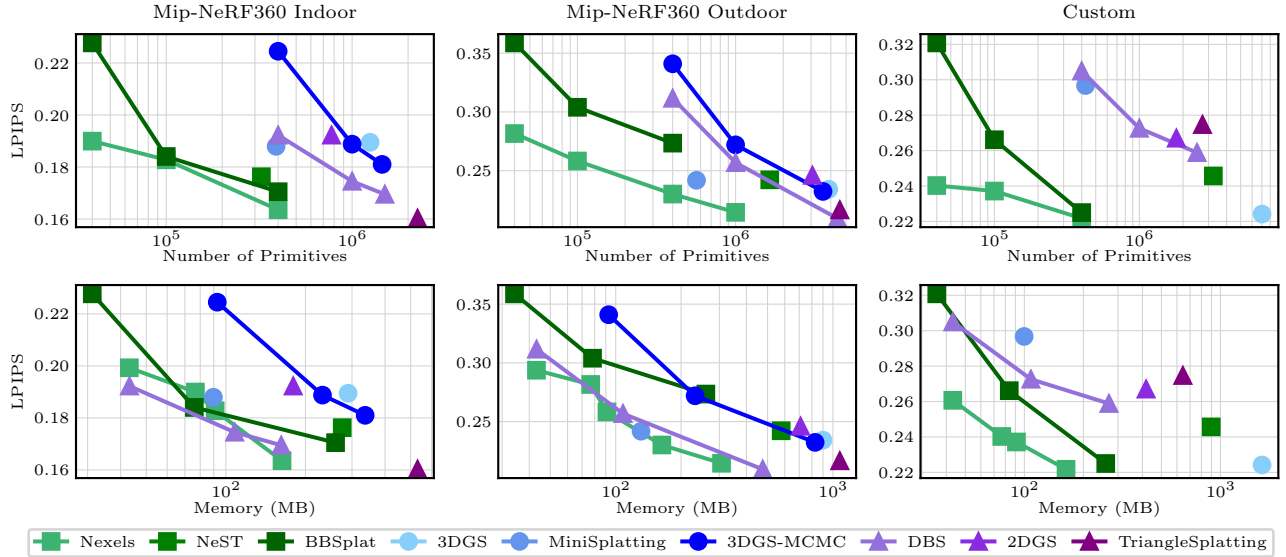
Figure 5. **Quality versus representation size.** We show how the perceptual quality of our representation varies depending on the number of primitives (top row) and memory (bottom row). We evaluate the LPIPS↓ across multiple settings for several methods and display separate graphs averaged over the four Mip-NeRF360 indoor scenes (left column), five Mip-NeRF360 outdoor scenes (middle column), and the four highly textured scenes of our custom dataset (right column). Nexels consistently has the lowest LPIPS for any number of primitives under 10^6, for all three of the datasets. When plotted against memory, nexels is always in the top three at any amount.

	Mip-NeRF360					Tanks & Temples					Custom				
	PSNR	SSIM	LPIPS	Memory	#Prim.	PSNR	SSIM	LPIPS	Memory	#Prim.	PSNR	SSIM	LPIPS	Memory	#Prim.
3DGS [22]	27.21	0.815	0.214	640MB	2.7M	23.80	0.853	0.169	348MB	1.5M	26.65	**0.858**	0.224	1630MB	7.0M
2DGS [20]	27.04	0.805	0.252	480MB	2.0M	23.15	0.831	0.212	200MB	0.8M	25.98	0.831	0.267	450MB	1.8M
DBS [30]	**28.10**	**0.829**	0.192	340MB	3.1M	**24.82**	**0.871**	0.144	260MB	2.0M	26.21	0.842	0.259	270MB	2.5M
TriSplat [16]	27.16	0.814	**0.191**	790MB	3.5M	23.14	0.857	**0.143**	400MB	2.0M	24.28	0.780	0.275	620MB	2.7M
NeST-Splat [59]	26.68	0.795	0.212	440MB	1.0M	23.02	0.824	0.181	260MB	0.5M	25.40	0.824	0.246	900MB	3.2M
BBSplat [49]	26.98	0.783	0.231	260MB	0.4M	23.83	0.854	0.147	260MB	0.4M	26.57	0.844	0.225	260MB	0.4M
Ours	27.35	0.802	0.201	160MB	0.4M	23.55	0.841	0.155	160MB	0.4M	**27.04**	0.851	**0.222**	160MB	0.4M

Table 2. **Quantitative results.** We compare nexels with 400K primitives to baselines at their maximum recommended number of primitives, corresponding to each baseline's rightmost point on Figure 5.

passes with autograd and use the Adam optimizer for all parameters [24, 39]. We implement custom CUDA kernels for our multiresolution hash-grid operations and rasterization passes. We use the NVIDIA tinycudann library for the fully-fused MLP pass [36]. For all scenes, we use a texture limit of just $K = 2$. The MLP has two hidden layers without bias terms, each with a width of 64, and ReLU activations. We use degree 3 spherical harmonics coefficients for both appearance modes, corresponding to $S = 16$ and an MLP output dimension of 48. For the hash-grid, $L = 16$ and $F = 2$. Unless otherwise stated, we use $T = 2^{20}$.

5.1. Datasets and Metrics

We perform our evaluation over fifteen real-world scenes across three datasets: all nine scenes from Mip-NeRF360 [3], the TRAIN and TRUCK scenes from the Tanks & Temples dataset [25], and four scenes which we collect involving high amounts of texture detail. Existing datasets seldom have close-up views despite these being important for photorealism. Fine appearance details within these scenes barely show up in their metrics. Our scenes contain high-frequency appearance, such as text and patterns, observed at a range of distances, including close-up shots. We describe the capture and preprocessing details in the supplementary. The full dataset will be released publicly.

We compare against point-based representations [16, 20, 22, 30] and textured representations [49, 59], which are state-of-the-art for novel view synthesis. We refer to the supplementary for experimental details for our baselines. In addition, we consider Gaussian splat variants which can produce sparser point sets than the original 3DGS [12, 23].

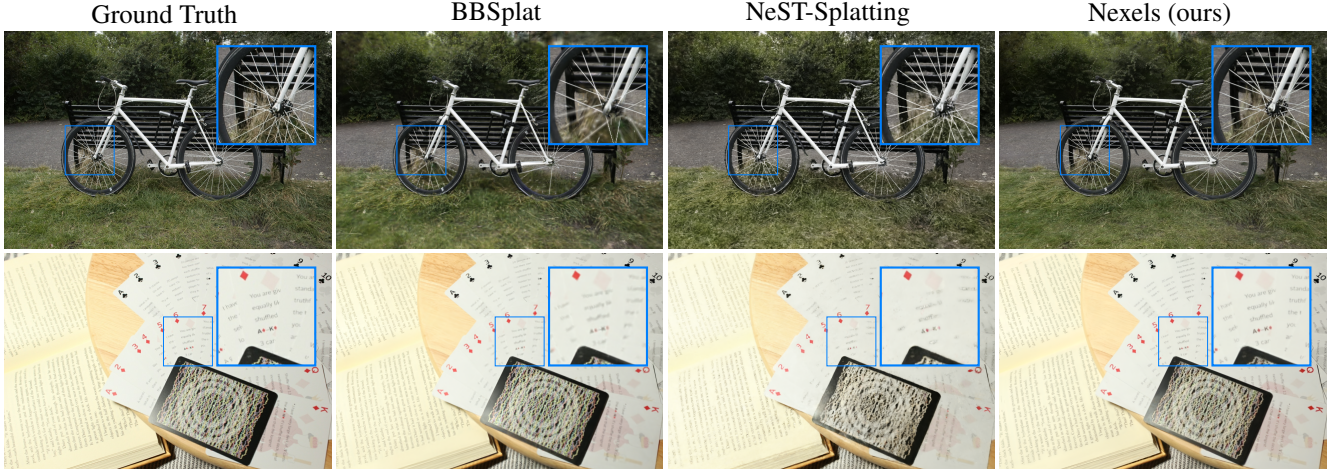We evaluate photometric quality with the standard

6

Figure 6. **Sparse geometries.** We show the BICYCLE and TABLE scenes reconstructed with 40K primitives across different textured representations. Only nexels are able to maintain a high quality in this extreme setting.
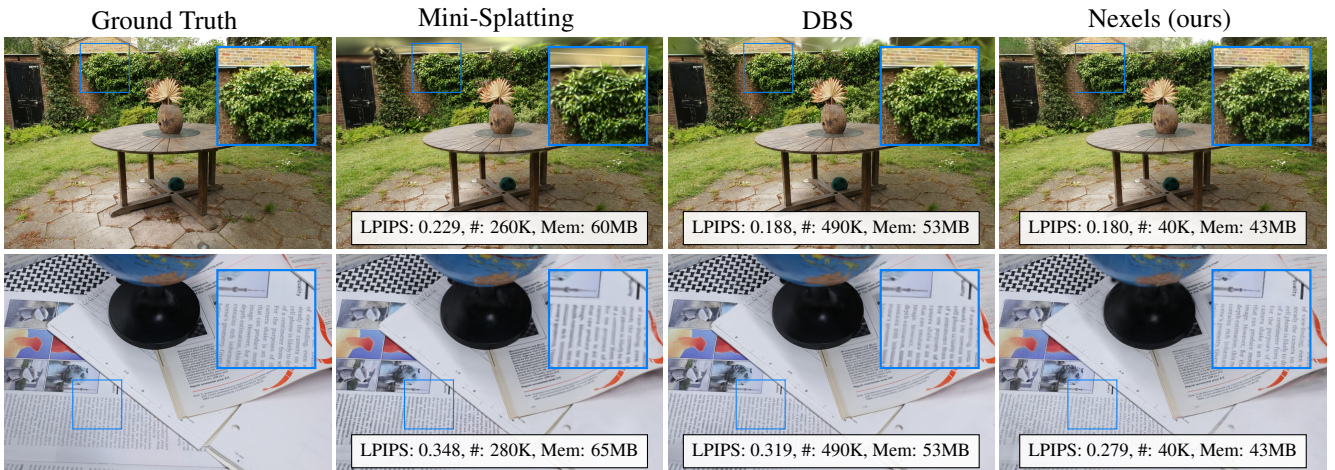

Figure 7. **Low memory comparisons.** We adjust the settings of baseline methods designed to support low memory representations to get a model around 50MB in size. We show the average LPIPS↓ across test views, number of primitives, and total memory for the GARDEN and TRIPOD scenes. With less memory and far fewer primitives, our nexels are able to capture fine appearance details and sparsely-observed background structures which point-based methods struggle with. We use $T = 2^{19}$ for the hash grid here.

PSNR, SSIM, and LPIPS metrics used in novel view synthesis and differentiable rendering [34, 58]. The LPIPS metric across all baselines is computed with the same normalization as the original 3DGS paper [22, 43]. All timings are performed on a single RTX 6000 Ada GPU using the CUDA implementation of each method's forward pass.

5.2. Evaluation

Novel-view synthesis. We compare all baselines at each of their maximum recommended settings in Tab. 2. On the Mip-NeRF360 dataset, nexels are competitive with point-based representations with only 400K primitives, achieving better PSNR than all methods except Deformable Beta Splatting (DBS [30]), which has over $7\times$ the number of

primitives and twice the memory. On our custom dataset, we achieve the highest PSNR and LPIPS and come second in SSIM to 3DGS at $10\times$ our memory. We achieve better scores than NeST-Splatting and BBSplat on the Mip-NeRF360 and custom datasets.

Visual quality under primitive and memory budgets. In Figure 5, we plot the perceptual quality (LPIPS) against the number of primitives and memory across three scene groups. For baselines which can specify a number of primitives, we include data points corresponding to appropriate subsets of $\{40K, 100K, 400K, 1M\}$ primitives. Details are included in the supplementary material.

Our representation exceeds virtually all other models at

Method	Indoor		Outdoor		Mean	
	FPS	Train.	FPS	Train.	FPS	Train.
NeST-Splat	27	4.4h	19	5.2h	23	4.9h
BBSplat	20	2.4h	20	2.1h	20	2.2h
Nexels	**40**	**1.5h**	**58**	**0.9h**	**50**	**1.1h**

Table 3. **Timing comparisons.** We show the average FPS of test-views across the Mip-NeRF360 scenes for the textured representations as well as training times.

	PSNR	SSIM	LPIPS	Mem.	FPS
No texture	27.20	0.787	0.261	100MB	**85**
No gamma	27.10	0.797	0.210	160MB	48
No prim. SH	27.00	0.797	0.205	**90MB**	47
No down.	27.03	0.797	0.211	160MB	49
Full	**27.35**	**0.802**	**0.201**	160MB	50

Table 4. **Ablations.** We show how dropping individual elements of nexels affects the photometric quality, memory, and FPS of the representation averaged across all Mip-NeRF360 scenes.

equal numbers of primitives. For example, 400K nexels obtain an LPIPS score of 0.164 across Mip-NeRF360 indoor scenes, while BBSplat and DBS obtain LPIPS scores of 0.171 and 0.192, respectively. Conversely, to reach a given LPIPS, nexels require far fewer primitives than other methods. The best result of 3DGS on the Mip-NeRF360 outdoor scenes is 0.234 LPIPS with 3.9 million Gaussians. We achieve 0.230 LPIPS with just 400K nexels. Similarly, 3DGS obtains 0.190 LPIPS on indoor scenes with 1.2 million Gaussians. We reach the same score at 40K nexels. Our method is also competitive when measured against memory. At 40MB on MipNeRF-360 outdoor scenes, we achieve 0.294 LPIPS while DBS achieves 0.312 LPIPS. We include qualitative results echoing these trends in Figures 6 and 7.

For the custom scenes with high amounts of texture, nexels are clearly better than non-textured representations. We achieve an LPIPS of 0.240 with just 40K nexels, outperforming DBS with 2.5 million splats, which obtains 0.259 LPIPS. BBSplat comes near us at 400K primitives but struggles at low primitive counts.

Training time and rendering speed. Table 3 compares training time and rendering speed for different texturing works. We separate between indoor and outdoor scenes as they have different image sizes. NeST-Splat averages five hours of training and only renders at 23 FPS across Mip-NeRF360 scenes. This is due to its texturing, where hash-grid lookups are done for each ray–primitive intersection, totaling thousands of random accesses per ray. BBSplat has a faster training time, but their use of alpha textures results in high overdraw, leading to a rendering speed of 20 FPS. In contrast, nexels train 4.2 and 1.9 times faster than NeST-Splatting and BBSplat, respectively, and render at 50 FPS.

5.3. Ablations

The results in Tab. 4 ablate individual components of our method to demonstrate their importance. We evaluate each individual experiment across the Mip-NeRF360 scenes trained with 400K primitives.

First, we verify that the main component, the neural texture, is useful. We find that all visual metrics, especially LPIPS, worsen without the neural texture. Qualitatively,

the non-textured renders lose background details. Removing the γ parameter and using 2D Gaussians as the kernel results in slight degradations overall as a Gaussian kernel leads to lower texture weights. Dropping the higher-order per-primitive spherical harmonics coefficients and only relying on the neural field for view-dependent effects is a possibility if memory is the main priority, as it removes 45 parameters for each primitive. The increase in LPIPS is very small, although there is a trade-off in PSNR. Finally, omitting the hash-grid downweighting factor in eq. (4) leads to noticeable aliasing artifacts in test views, as the fine levels overfit at the intersections seen from training poses.

5.4. Limitations

The speed of the Instant-NGP model is reliant on tensor core operations, which are only available on high-end GPUs. While nexels are real-time on the machinery we use, real-time rendering on mobile GPUs or low-end laptop GPUs is currently not possible. In addition, nexels produce noise in unseen regions, which is likely less appealing for end users than the blurry artifacts of Gaussian splats. Our work also does not support the full complexity of the imaging process, including motion blur and depth-of-field blur. These effects are particularly relevant for differentiable rendering of fine details, as even small deviations from the expected pinhole camera model can damage reconstructions.

6. Conclusion

We present the first representation for novel view synthesis that removes the reliance on dense primitives while achieving real-time rendering at high quality. Our results show that far fewer geometries are needed to model scene-level datasets than recent research would suggest. Future work could incorporate level-of-detail structures to model city-scale scenes. Sparse surfaces also work well for ray-tracing applications. Finally, the top-K scheme can be extended to support textures combining the neural field and per-primitive features. In conclusion, nexels break new ground in the exploration between volumetric and surface representations, capturing high-frequency appearance at a low memory and compute cost.

References

[1] Tomas Akenine-Möller, Eric Haines, Naty Hoffman, Angelo Pesce, Michał Iwanicki, and Sébastien Hillaire. *Real-Time Rendering*. AK Peters/CRC Press, Boca Raton, Florida, USA, 4 edition, 2018. 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 *IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, pages 5835–5844, Montréal, Can., 2021. Inst. Electr. Electron. Eng. (IEEE). 1, 2

[3] Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-NeRF 360: Unbounded anti-aliased neural radiance fields. In *IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 5460–5469, New Orleans, LA, USA, 2022. Inst. Electr. Electron. Eng. (IEEE). 2, 6, 1

[4] Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Zip-NeRF: Anti-aliased grid-based neural radiance fields. In *IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, pages 19640–19648, Paris, Fr., 2023. Inst. Electr. Electron. Eng. (IEEE). 4, 5, 1

[5] Louis Bavoil, Steven P. Callahan, Aaron Lefohn, João L. D. Comba, and Cláudio T. Silva. Multi-fragment effects on the GPU using thek-buffer. In *ACM Symp. Interact. 3D Graph. Games*, pages 97–104, Seattle, Washington, 2007. ACM. 5

[6] Ang Cao and Justin Johnson. HexPlane: A fast representation for dynamic scenes. In *IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 130–141, Vancouver, Can., 2023. IEEE. 2

[7] Brian Chao, Hung-Yu Tseng, Lorenzo Porzi, Chen Gao, Tuotuo Li, Qinbo Li, Ayush Saraf, Jia-Bin Huang, Johannes Kopf, Gordon Wetzstein, and Changil Kim. Textured Gaussians for enhanced 3D scene appearance modeling. In *IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 8964–8974, Nashville, TN, USA, 2025. IEEE. 2, 3

[8] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. TensoRF: Tensorial radiance fields. In *Eur. Conf. Comput. Vis. (ECCV)*, pages 333–350, Tel Aviv, Israël, 2022. Springer Nat. Switz. 2

[9] Hanyu Chen, Bailey Miller, and Ioannis Gkioulekas. 3D reconstruction with fast dipole sums. *ACM Trans. Graph.*, 43 (6):1–19, 2024. 2

[10] Zhang Chen, Zhong Li, Liangchen Song, Lele Chen, Jingyi Yu, Junsong Yuan, and Yi Xu. NeuRBF: A neural fields representation with adaptive radial basis functions. In *IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, pages 4159–4171, Paris, Fr., 2023. IEEE. 2

[11] Pinxuan Dai, Jiamin Xu, Wenxiang Xie, Xinguo Liu, Huamin Wang, and Weiwei Xu. High-quality surface reconstruction using Gaussian surfels. In *ACM SIGGRAPH Conf. Pap.*, pages 1–11, Denver, CO, USA, 2024. ACM. 2

[12] Guangchi Fang and Bing Wang. Mini-splatting: Representing scenes with a constrained number of Gaussians. In *Eur. Conf. Comput. Vis. (ECCV)*, pages 165–181. Springer Nat. Switz., 2024. 6, 4

[13] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 5491–5500, New Orleans, LA, USA, 2022. Inst. Electr. Electron. Eng. (IEEE). 2

[14] Yuan-Chen Guo, Yan-Pei Cao, Chen Wang, Yu He, Ying Shan, and Song-Hai Zhang. VMesh: Hybrid volume-mesh representation for efficient view synthesis. In *ACM SIGGRAPH Asia Conf. Pap.*, pages 1–11, Sydney, New South Wales, Aust., 2023. ACM. 2

[15] Abdullah Hamdi, Luke Melas-Kyriazi, Jinjie Mai, Guocheng Qian, Ruoshi Liu, Carl Vondrick, Bernard Ghanem, and Andrea Vedaldi. GES: Generalized exponential splatting for efficient radiance field rendering. In *IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 19812–19822, Seattle, WA, USA, 2024. Inst. Electr. Electron. Eng. (IEEE). 3, 4

[16] Jan Held, Renaud Vandeghen, Adrien Deliège, Daniel Hamdi, Abdullah Rebain, Silvio Giancola, Anthony Cioppa, Andrea Vedaldi, Bernard Ghanem, Andrea Tagliasacchi, and Marc Van Droogenbroeck. Triangle splatting for real-time radiance field rendering. *arXiv*, abs/2505.19175, 2025. 6

[17] Jan Held, Renaud Vandeghen, Abdullah Hamdi, Adrien Deliège, Anthony Cioppa, Silvio Giancola, Andrea Vedaldi, Bernard Ghanem, and Marc Van Droogenbroeck. 3D convex splatting: Radiance field rendering with 3D smooth convexes. In *IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 21360–21369, Nashville, TN, USA, 2025. Inst. Electr. Electron. Eng. (IEEE).

[18] Jan Held, Renaud Vandeghen, Sanghyun Son, Daniel Rebain, Matheus Gadelha, Yi Zhou, Ming C. Lin, Marc Van Droogenbroeck, and Andrea Tagliasacchi. Triangle splatting+: Differentiable rendering with opaque triangles. *arXiv*, abs/2509.25122, 2025. 3

[19] Jan Held, Renaud Vandeghen, Adrien Deliège, Daniel Hamdi, Abdullah Rebain, Silvio Giancola, Anthony Cioppa, Andrea Vedaldi, Bernard Ghanem, Andrea Tagliasacchi, and Marc Van Droogenbroeck. Triangle splatting for real-time radiance field rendering. In *Int. Conf. 3D Vis. (3DV)*, pages 1–10, Vancouver, Can., 2026. 2

[20] Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2D Gaussian splatting for geometrically accurate radiance fields. In *ACM SIGGRAPH Conf. Pap.*, pages 1–11, Denver, CO, USA, 2024. ACM. 2, 3, 4, 6

[21] John Hughes, Andries van Dam, Morgan McGuire, David Sklar, James D. Foley, Steven Feiner, and Kurt Akeley. *Computer Graphics: Principles and Practice*. Addison-Wesley, 3 edition, 2014. 2

[22] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkuehler, and George Drettakis. 3D Gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):1–14, 2023. 1, 2, 5, 6, 7

[23] Shakiba Kheradmand, Daniel Rebain, Gopal Sharma, Wei-wei Sun, Jeff Tseng, Hossam Isack, Abhishek Kar, Andrea Tagliasacchi, and Kwang Moo Yi. 3D Gaussian splatting as Markov chain Monte Carlo. In *Adv. Neural Inf. Process. Syst. (NeurIPS)*, pages 80965–80986, Vancouver, Can., 2024. Curran Assoc. Inc. 5, 6, 2

[24] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv*, abs/1412.6980, 2014. 6

[25] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: benchmarking large-scale scene reconstruction. *ACM Trans. Graph.*, 36(4):1–13, 2017. 6

[26] Samuli Laine, Janne Hellsten, Tero Karras, Yeongho Seol, Jaakko Lehtinen, and Timo Aila. Modular primitives for high-performance differentiable rendering. *ACM Trans. Graph.*, 39(6):1–14, 2020. 2

[27] Jingtao Li, Jian Zhou, Yan Xiong, Xing Chen, and Chaitali Chakrabarti. An adjustable farthest point sampling method for approximately-sorted point cloud data. In *IEEE Work. Signal Process. Syst. (sips)*, pages 1–6, Rennes, France, 2022. IEEE. 5, 1, 2, 4

[28] Ruilong Li, Hang Gao, Matthew Tancik, and Angjoo Kanazawa. NerfAcc: Efficient sampling accelerates NeRFs. In *IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, pages 18491–18500, Paris, Fr., 2023. IEEE. 2

[29] Tzu-Mao Li, Miika Aittala, Frédo Durand, and Jaakko Lehtinen. Differentiable Monte Carlo ray tracing through edge sampling. *ACM Trans. Graph.*, 37(6):1–11, 2018. 2

[30] Rong Liu, Dylan Sun, Meida Chen, Yue Wang, and Andrew Feng. Deformable beta splatting. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Paper*, pages 1–11. ACM, 2025. 3, 4, 6, 7

[31] Shichen Liu, Weikai Chen, Tianye Li, and Hao Li. Soft rasterizer: A differentiable renderer for image-based 3D reasoning. In *IEEE Int. Conf. Comput. Vis. (ICCV)*, pages 7707–7716, Seoul, South Korea, 2019. Inst. Electr. Electron. Eng. (IEEE). 2

[32] Tao Lu, Mulin Yu, Linning Xu, Yuanbo Xiangli, Limin Wang, Dahua Lin, and Bo Dai. Scaffold-GS: Structured 3D Gaussians for view-adaptive rendering. In *IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 20654–20664, Seattle, WA, USA, 2024. IEEE. 3

[33] Julien N. P. Martel, David B. Lindell, Connor Z. Lin, Eric R. Chan, Marco Monteiro, and Gordon Wetzstein. Acorn: : adaptive coordinate networks for neural scene representation. *ACM Trans. Graph.*, 40(4):1–13, 2021. 2

[34] 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. *Commun. ACM*, 65(1):99–106, 2021. 2, 7

[35] Nicolas Moenne-Loccoz, Ashkan Mirzaei, Or Perel, Riccardo de Lutio, Janick Martinez Esturo, Gavriel State, Sanja Fidler, Nicholas Sharp, and Zan Gojcic. 3D Gaussian ray tracing: Fast tracing of particle scenes. *ACM Trans. Graph.*, 43(6):1–19, 2024. 3

[36] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a mul-tiresolution hash encoding. *ACM Trans. Graph.*, 41(4):1–15, 2022. 2, 3, 4, 6, 1

[37] Baptiste Nicolet, Alec Jacobson, and Wenzel Jakob. Large steps in inverse rendering of geometry. *ACM Trans. Graph.*, 40(6):1–13, 2021. 2

[38] Panagiotis Papantonakis, Georgios Kopanas, Fredo Durand, and George Drettakis. Content-aware texturing for gaussian splatting. *arXiv preprint arXiv:2512.02621*, 2025. 2, 3

[39] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Adv. Neural Inf. Process. Syst. (NeurIPS)*, pages 8026–8037, Vancouver, Can., 2019. Curran Assoc. Inc. 6

[40] Christian Reiser, Stephan Garbin, Pratul Srinivasan, Dor Verbin, Richard Szeliski, Ben Mildenhall, Jonathan Barron, Peter Hedman, and Andreas Geiger. Binary opacity grids: Capturing fine geometric detail for mesh-based view synthesis. *ACM Trans. Graph.*, 43(4):1–14, 2024. 2

[41] Kerui Ren, Lihan Jiang, Tao Lu, Mulin Yu, Linning Xu, Zhangkai Ni, and Bo Dai. Octree-GS: Towards consistent real-time rendering with LOD-structured 3D Gaussians. *IEEE Trans. Pattern Anal. Mach. Intell.*, pages 1–15, 2025. 3

[42] Victor Rong, Jingxiang Chen, Sherwin Bahmani, Kiriakos N. Kutulakos, and David B. Lindell. GStex: Per-primitive texturing of 2D Gaussian splatting for decoupled appearance and geometry modeling. In *IEEE/CVF Winter Conf. Appl. Comput. Vis. (WACV)*, pages 3508–3518, Tucson, AZ, USA, 2025. IEEE. 2, 3

[43] Samuel Rota Bulò, Lorenzo Porzi, and Peter Kontschieder. Revising densification in Gaussian splatting. In *Eur. Conf. Comput. Vis. (ECCV)*, pages 347–362. Springer Nat. Switz., 2024. 5, 7, 2

[44] Johannes L. Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 4104–4113, Las Vegas, NV, USA, 2016. Inst. Electr. Electron. Eng. (IEEE). 5

[45] Gopal Sharma, Daniel Rebain, Kwang Moo Yi, and Andrea Tagliasacchi. Volumetric rendering with baked quadrature fields. In *Eur. Conf. Comput. Vis. (ECCV)*, pages 275–292. Springer Nat. Switz., 2024. 2

[46] Nicholas Sharp and Alec Jacobson. Spelunking the deep. *ACM Trans. Graph.*, 41(4):1–16, 2022. 2

[47] Christian Sigg, Tim Weyrich, Mario Botsch, and Markus Gross. GPU-based ray-casting of quadratic surfaces. In *Symposium on Point-Based Graphics*, Boston, MA, USA, 2006. Eurographics Assoc. 3

[48] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 5449–5459, New Orleans, LA, USA, 2022. Inst. Electr. Electron. Eng. (IEEE). 2

[49] David Svitov, Pietro Morerio, Lourdes Agapito, and Alessio Del Bue. Billboard splatting (BBSplat): Learnable textured primitives for novel view synthesis. In *IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, pages 25029–25039, Honolulu, HI, USA, 2025. IEEE. 1, 2, 3, 5, 6

[50] Towaki Takikawa, Joey Litalien, Kangxue Yin, Karsten Kreis, Charles Loop, Derek Nowrouzezahrai, Alec Jacobson, Morgan McGuire, and Sanja Fidler. Neural geometric level of detail: Real-time rendering with implicit 3D shapes. In *IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 11353–11362, Nashville, TN, USA, 2021. IEEE. 2

[51] Zian Wang, Tianchang Shen, Jun Gao, Shengyu Huang, Jacob Munkberg, Jon Hasselgren, Zan Gojcic, Wenzheng Chen, and Sanja Fidler. Neural fields meet explicit geometric representations for inverse rendering of urban scenes. In *IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 8370–8380, Vancouver, Can., 2023. IEEE. 2

[52] Zian Wang, Tianchang Shen, Merlin Nimier-David, Nicholas Sharp, Jun Gao, Alexander Keller, Sanja Fidler, Thomas Müller, and Zan Gojcic. Adaptive shells for efficient neural radiance field rendering. *ACM Trans. Graph.*, 42(6):1–15, 2023. 2

[53] Qiangeng Xu, Zexiang Xu, Julien Philip, Sai Bi, Zhixin Shu, Kalyan Sunkavalli, and Ulrich Neumann. Point-NeRF: Point-based neural radiance fields. In *IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 5428–5438, New Orleans, LA, USA, 2022. IEEE. 2

[54] Tian-Xing Xu, Wenbo Hu, Yu-Kun Lai, Ying Shan, and Song-Hai Zhang. Texture-GS: Disentangling the geometry and texture for 3D Gaussian splatting editing. In *Eur. Conf. Comput. Vis. (ECCV)*, pages 37–53. Springer Nat. Switz., 2024. 3

[55] Lior Yariv, Peter Hedman, Christian Reiser, Dor Verbin, Pratul P. Srinivasan, Richard Szeliski, Jonathan T. Barron, and Ben Mildenhall. BakedSDF: Meshing neural SDFs for real-time view synthesis. In *ACM SIGGRAPH Conf. Proc.*, pages 1–9, Los Angeles, CA, USA, 2023. ACM. 2

[56] Keyang Ye, Tianjia Shao, and Kun Zhou. When Gaussian meets surfel: Ultra-fast high-fidelity radiance field rendering. *ACM Trans. Graph.*, 44(4):1–15, 2025. 2

[57] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. PlenOctrees for real-time rendering of neural radiance fields. In *IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, pages 5732–5741, Montréal, Can., 2021. IEEE. 2

[58] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 586–595, Salt Lake City, UT, USA, 2018. IEEE. 7, 5

[59] Xin Zhang, Anpei Chen, Jincheng Xiong, Pinxuan Dai, Yujun Shen, and Weiwei Xu. Neural shell texture splatting: More details and fewer primitives. In *IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, pages 25229–25238, Honolulu, HI, USA, 2025. IEEE. 1, 2, 3, 6, 4

Nexels: Neurally-Textured Surfels
for Real-Time Novel View Synthesis with Sparse Geometries

Supplementary Material

7. Overview

We include additional implementation details, dataset details, experiment details, and discussion below.

8. Implementation Details

8.1. Instant-NGP

Architecture. The original Instant-NGP implementation assumes inputs $x \in [0,1]^d$. Later works apply it to large-scale scenes by contracting space to this hypercube [3, 4, 59]. We instead make minor modifications to the hash-grid in order to accept unbounded inputs. First, we use a hash-grid, even for the coarse levels. Second, the original hash function used by Müller et al. [36] has structured collisions when evaluated over both positive and negative inputs. We modify the hash to accept negative integers by applying the function

$$\texttt{MapPositive}(x) = \begin{cases} 2x - 1 & x > 0, \\ -2x & \text{otherwise,} \end{cases}$$

to the hash function inputs.

Anti-aliasing. As mentioned in the main text, we additionally perform anti-aliasing on the hash-grid using the down-weighting strategy with grid weight decay proposed in ZipNeRF [4]. We describe how we derive our down-weighting factor, which differs slightly from that of Barron et al. [4].

Let t^* be the depth along the ray and f to be the focal length. We model the projected pixel footprint as a 2D isotropic normal distribution $\mathcal{N}$ in world-space, with mean $x^* := \mathbf{o} + t^*\mathbf{d}$ and standard deviation $\frac{t^*}{f}$, that is parallel to the image plane. The integrated texture which we seek to approximate is $\mathbb{E}_{x \sim \mathcal{N}}[\mathcal{F}(x)]$.

Estimating the MLP as a linear function, we have

$$\mathbb{E}_{x \sim \mathcal{N}}[\mathcal{F}(x)] \approx \mathcal{G}\left(\mathbb{E}_{x \sim \mathcal{N}}[\mathcal{H}(x)]\right).$$

The hash-grid at level ℓ, $\mathcal{H}_\ell$, is queried by hashing $s_\ell x$ and using trilinear interpolation. For a coarse but efficient estimate of $\mathbb{E}_{x \sim \mathcal{N}}[\mathcal{H}_\ell(s_\ell x)]$, we only query $\mathcal{H}_\ell$ at $s_\ell x^*$ and approximate the rest. More formally, we model

$$\mathcal{H}_\ell(s_\ell x^* + \epsilon) \approx w(\epsilon)\mathcal{H}_\ell(s_\ell x^*) + (1 - w(\epsilon))\mathcal{H}_\ell(s_\ell x^* + \epsilon),$$

where $w(\epsilon)$ indicates a unit cube, i.e. $w(\epsilon) = 1$ if $||\epsilon||_\infty \leq \frac{1}{2}$ and $w(\epsilon) = 0$ otherwise. We assume that the unknown values are zero in expectation. Hence,

$$\mathbb{E}_{x \sim \mathcal{N}}[\mathcal{H}_\ell(s_\ell x)] \approx \mathbb{E}_{x \sim \mathcal{N}}[w(s_\ell(x - x^*))\mathcal{H}_\ell(s_\ell x^*)]$$

$$\approx \text{erf}\left(\frac{f}{2\sqrt{2}s_\ell t^*}\right)^2 \mathcal{H}_\ell(s_\ell x^*).$$

Using the approximation $\text{erf}(x)^2 \approx 1 - \exp\left(\frac{-4x^2}{\pi}\right)$ yields our down-weighting factor

$$\Delta_\ell := 1 - \exp\left(-\frac{1}{2\pi}\left(\frac{f}{s_\ell t^*}\right)^2\right).$$

The above derivation makes several very crude approximations. It's likely that many down-weighting factors are adequate. More high-powered anti-aliasing techniques could be used, but we found that this simple down-weighting was already enough to remove noticeable artifacts.

Note that the MLP in the tinycudann library does not use bias terms. As we use ReLU activations, a natural consequence is that $\mathcal{G}(\vec{0}) = \vec{0}$. We choose to preserve this property as our derivation approximated $\mathcal{G}$ to be linear, but we observe that the neural texture is mostly grey in the background of large scenes, particularly the TRAIN scene of the Tanks & Temples dataset.

Other details. Unlike prior work in which the sample positions are assumed to be non-differentiable [36, 59], we compute the positional gradient from the bilinear interpolation into each hash-grid level.

Both the hash-grid and MLP are computed using 16-bit floating point numbers. The latter's quantization is particularly important in order to leverage tensor core operations on recent NVIDIA GPUs, which greatly accelerates the MLP pass. The gradients are accumulated within 32-bit floating point tensors to maintain accuracy during optimization. We use the Instant-NGP codebase for the MLP pass and implement our own hash-grid following the design specified above. When performing the backward pass, we use the same default loss scale of 128 as Instant-NGP to reduce underflow. We report the nexels storage memory using two bytes for each Instant-NGP parameter and four bytes for all other parameters.

8.2. Adaptive Density Control

Initialization. Given N points $x_0, \ldots, x_{N-1}$, farthest point sampling iteratively samples the point which is farthest from the ones already sampled [27]. In other words,

it returns an indexing $I_0, I_1, \ldots, I_{N-1}$, such that for any $0 < i \leq N - 1$,

$$I_i := \operatorname*{arg\,max}_{k \notin \{I_0, \ldots, I_{i-1}\}} \min_{0 \leq j < i} d\left(x_k, x_{I_j}\right).$$

When initializing our model, we use farthest point sampling to sample the positions. For the scale, the nearest-neighbors initialization suggested by Kerbl et al. [22] is not ideal for our representation as it leads to small primitives. Instead, we run nearest-neighbors across prefixes of the farthest point sampling order so that the initial scale for point i with position x_{I_i} is approximately $\min_{0 \leq j < i} d\left(x_{I_i}, x_{I_j}\right)$. This implicitly leads to a hierarchy where points sampled early are larger while points sampled later are smaller. These smaller primitives can potentially be removed if the optimization finds that the larger primitives can replace them with texture alone.

Densification. In the backwards pass of each iteration, we return the blended L_1 errors as defined in the work of Rota Bulò et al. [43]. These blended errors are a per-primitive heuristic quantity where a higher error suggests that the primitive should be split. We accumulate these errors across iterations by adding them, rather than performing primitive-wise maximums. Inspired by 3DGS-MCMC [23], we normalize the blended errors and use them as probability weights to sample 5% of the current primitives for splitting.

Our splitting operation aims to preserve the textured surface. Say we are splitting surfel i. We will assume $\sigma_{i,1} \geq \sigma_{i,2}$ for convenience. The surfel is then split into two primitives both with scales $(\frac{\sigma_{i,1}}{2}, \sigma_{i,2})$ and with means $\mu_i - \frac{\sigma_{i,1}}{2}\mathbf{v}_{i,1}$ and $\mu_i + \frac{\sigma_{i,1}}{2}\mathbf{v}_{i,1}$, respectively. All other properties of the new primitives, including the γ values, are copied from the original one. For $\gamma_i \to \infty$, this splits the quad perfectly in half along its long axis.

9. Dataset Details

We capture four scenes, GRAFFITI, GROCERY, TABLE, and TRIPOD. All images were captured using a Canon EOS Rebel T7 equipped with an EF-S 18-55mm f/3.5-5.6 IS II zoom lens set to a focal length of approximately 35 mm.

To ensure photometric consistency, we configured camera settings by capturing preliminary reference shots from multiple angles at each scene. We then locked these settings for the capture of each scene, fixing ISO, white balance, zoom level, and exposure per scene. Auto-focus was enabled. These scenes were selected specifically to evaluate reconstruction performance on high-frequency textures. A similar dataset was captured by Chao et al. [7], but the data was not made public for licensing reasons.

For the *indoor* scenes (GROCERY, TABLE, TRIPOD), we captured 360-degree concentric trajectories centered on the

object of interest. For the *outdoor* scene (GRAFFITI), which lacks a singular central subject, we densely captured the region of interest. It took approximately 30 minutes to capture up to 300 images for each scene. After capturing, we manually removed images with any noticeable motion blur. We include brief descriptions of each scene below. The dataset will be made public.

- GRAFFITI (254 images). This outdoor scene depicts an alleyway where a concrete driveway separates a fence from a graffiti-covered building.
- GROCERY (278 images). Captured indoors, this scene centers on a 0.5×1 meter dining table covered in kitchen rags. On top of the rags, the table features a variety of grocery items with text-heavy labels as well as an open book.
- TABLE (165 images). This scene was taken inside a single room, with the camera focusing on a small table around 1 meter in length. Two open books, a newspaper, and a deck of cards are laid out on the table, along with miscellaneous other objects.
- TRIPOD (210 images). Arranged inside a small room, this scene consists of magazines, markers, and various boxes surrounding a large poster propped up on a tripod.

10. Experiment Details

10.1. Baseline Details

We evaluated the baselines at various settings in order to obtain data points at a varying amount of primitives and memory. For 3DGS [22], 2DGS [20], and TriSplat [19], there is no fine control over the number of primitives. We use their default densification settings with no notable failures. For methods that provide separate hyperparameter sets for *outdoor* and *indoor* scenes, we used the *indoor* settings for our custom scenes. Additional details are provided below.

3DGS-MCMC [23]. The original paper uses the bicubic filter from the PIL library to downscale the images, specifically through the `resolution` flag. This produces smoother images than the ImageMagick downscaling which was used by Barron et al. [2] on the Mip-NeRF360 dataset. As a result, scripts yield slightly better photometric results when using the `resolution` flag for downscaling rather than using the `images` flag to specify the image directories downscaled by ImageMagick. For fairness, we re-evaluated 3DGS-MCMC using the same downscaling strategy as in the original 3D Gaussian Splatting work. When training for a low amount of splats, we initialize using farthest point sampling with the same number of points as our initialization [27].

BBSplat [49]. For BBSplat, we found that their scripts are run at a different resolution than the majority of Gaussian

	Mip-NeRF360 Indoor					Mip-NeRF360 Outdoor					Custom				
	PSNR	SSIM	LPIPS	Memory	#Prim.	PSNR	SSIM	LPIPS	Memory	#Prim.	PSNR	SSIM	LPIPS	Memory	#Prim.
3DGS	30.40	0.920	0.190	290MB	1.25M	24.64	0.731	0.234	900MB	3.88M	26.65	0.858	0.224	1630MB	7.03M
2DGS	30.41	0.917	0.192	180MB	0.78M	24.34	0.725	0.246	710MB	3.06M	25.98	0.831	0.267	420MB	1.80M
MiniSplatting	30.43	0.922	0.188	90MB	0.39M	24.72	0.741	0.242	130MB	0.57M	24.88	0.803	0.297	100MB	0.42M
TriSplat	30.80	0.928	0.160	530MB	2.25M	24.24	0.722	0.217	1080MB	4.56M	24.28	0.780	0.275	640MB	2.73M
NeST-Splat	30.35	0.909	0.176	280MB	0.33M	23.74	0.703	0.242	580MB	1.64M	25.39	0.824	0.246	900MB	3.24M
3DGS-MCMC (400K)	30.13	0.907	0.225	90MB	0.40M	24.03	0.672	0.341	90MB	0.40M	25.61	0.807	0.327	93MB	0.40M
3DGS-MCMC (1M)	31.16	0.926	0.189	230MB	1.00M	24.74	0.722	0.272	230MB	1.00M	25.96	0.823	0.295	232MB	1.00M
3DGS-MCMC	31.34	0.929	0.181	340MB	1.45M	25.04	0.744	0.232	830MB	3.57M	-	-	-	-	-
DBS (400K)	31.25	0.924	0.192	40MB	0.40M	24.11	0.692	0.312	40MB	0.40M	26.39	0.828	0.305	40MB	0.40M
DBS (1M)	31.98	0.933	0.174	110MB	1.00M	24.56	0.727	0.257	110MB	1.00M	26.72	0.844	0.273	110MB	1.00M
DBS	32.18	0.935	0.169	160MB	1.50M	24.83	0.745	0.209	470MB	4.40M	26.22	0.842	0.259	270MB	2.50M
BBSplat (40K)	28.85	0.886	0.228	30MB	0.04M	22.39	0.592	0.359	30MB	0.04M	24.98	0.794	0.321	40MB	0.04M
BBSplat (100K)	30.34	0.918	0.184	80MB	0.10M	23.17	0.647	0.304	80MB	0.10M	25.84	0.829	0.266	80MB	0.10M
BBSplat (400K)	31.16	0.926	0.170	260MB	0.40M	23.71	0.675	0.273	260MB	0.40M	26.57	0.844	0.225	260MB	0.40M
Nexels (40K, $T = 2^{19}$)	29.19	0.894	0.199	40MB	0.04M	23.26	0.641	0.294	40MB	0.04M	26.20	0.825	0.261	40MB	0.04M
Nexels (40K)	29.38	0.900	0.190	80MB	0.04M	23.25	0.647	0.282	80MB	0.04M	26.50	0.836	0.240	80MB	0.04M
Nexels (100K)	29.94	0.906	0.183	90MB	0.10M	23.81	0.675	0.258	90MB	0.10M	26.62	0.844	0.237	90MB	0.10M
Nexels (400K)	30.94	0.921	0.164	160MB	0.40M	24.48	0.707	0.230	160MB	0.40M	27.04	0.851	0.222	160MB	0.40M
Nexels (1M)	31.32	0.926	0.155	310MB	1.00M	24.77	0.722	0.215	310MB	1.00M	-	-	-	-	-

Table 5. **Quantitative results.** We show full results of the data points from Figure 5 of the main paper, as well as additional experiment configurations.

	Mip-NeRF360 Indoor				Mip-NeRF360 Outdoor					Tanks & Temples		Custom			
	Bns	Cnt	Ktc	Rm	Bcc	Flw	Grd	Stm	Trh	Trn	Trc	Grf	Grc	Tbl	Trp
Nexels (40K, $T = 2^{19}$)	29.67	27.44	29.41	30.23	23.36	19.85	25.48	25.08	22.54	-	-	24.86	26.93	27.88	25.14
Nexels (40K)	29.98	27.61	29.62	30.32	23.44	19.79	25.59	25.11	22.33	-	-	25.38	27.16	27.95	25.50
Nexels (100K)	30.81	28.13	30.32	30.51	24.05	20.38	26.23	25.88	22.53	-	-	25.50	27.51	27.73	25.74
Nexels (400K)	31.92	28.95	31.22	31.66	24.72	21.19	26.96	26.56	22.99	21.50	25.61	25.98	28.07	27.98	26.12
Nexels (1M)	32.47	29.26	31.59	31.97	25.04	21.48	27.36	27.01	22.97	-	-	-	-	-	-
Nexels (40K, $T = 2^{19}$)	0.907	0.875	0.895	0.899	0.645	0.488	0.784	0.688	0.600	-	-	0.743	0.853	0.901	0.802
Nexels (40K)	0.913	0.880	0.901	0.903	0.657	0.492	0.793	0.691	0.600	-	-	0.767	0.857	0.905	0.815
Nexels (100K)	0.923	0.892	0.910	0.899	0.690	0.530	0.817	0.728	0.607	-	-	0.776	0.864	0.904	0.831
Nexels (400K)	0.935	0.905	0.921	0.922	0.730	0.577	0.842	0.764	0.624	0.802	0.880	0.789	0.878	0.906	0.829
Nexels (1M)	0.942	0.911	0.927	0.925	0.749	0.595	0.855	0.781	0.629	-	-	-	-	-	-
Nexels (40K, $T = 2^{19}$)	0.211	0.214	0.154	0.218	0.304	0.376	0.178	0.290	0.322	-	-	0.338	0.264	0.164	0.277
Nexels (40K)	0.199	0.205	0.146	0.210	0.288	0.364	0.164	0.282	0.309	-	-	0.302	0.248	0.151	0.261
Nexels (100K)	0.188	0.192	0.136	0.215	0.261	0.334	0.145	0.249	0.302	-	-	0.318	0.234	0.146	0.251
Nexels (400K)	0.170	0.176	0.123	0.186	0.227	0.298	0.124	0.216	0.286	0.197	0.112	0.274	0.222	0.146	0.245
Nexels (1M)	0.161	0.167	0.114	0.179	0.208	0.281	0.110	0.197	0.277	-	-	-	-	-	-

Table 6. **Per-scene results.** We provide per-scene PSNR, SSIM, and LPIPS scores for our nexels at various settings. Scene names are abbreviated to their first three consonants.

splatting works. We re-ran the baseline using the same resolution as in 3D Gaussian Splatting. Additionally, we set a maximum primitive cap of 400K and limited the number of SfM points for initialization to 150K. For all scenes, we use their sky-box initialization that adds another 10K points in the background. For BBSplat, we were unable to reproduce the results for the STUMP scene when run at the standard resolution, as it achieved a significantly lower PSNR. For this scene only, we instead ran the script at their resolution and rescaled the renders to our resolution with Lanczos resampling. We evaluate the storage memory of BBSplat after

it performs dictionary-based compression on the textures.

DBS [30]. Like 3DGS-MCMC, the results in the original DBS paper were obtained using downsampling with the `resolution` flag. We re-ran the method using Mip-NeRF360's downscaled images. Furthermore, the original DBS uses a scene-specific number of iterations that explicitly checks the test set as a stopping condition. For fairness, we use the same number of iterations for all scenes (i.e. 30K). On our custom scenes, we set the primitive cap for DBS to 3 million primitives for GRAFFITI, GROCERY,
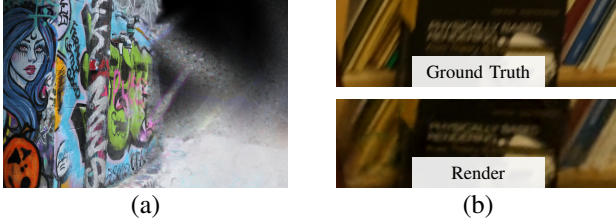
Figure 8. **Artifacts.** (a) We show a render of the GRAFFITI scene from a view outside of both the training and test sets. The neural field produces noise in unseen regions. (b) We show a zoom-in of the ground truth (top) and render (bottom) of the ROOM scene. Even small amounts of motion blur are destructive for fine details, and our reconstruction cannot properly capture the text.

K	PSNR	SSIM	LPIPS
0	24.81	0.663	0.415
1	25.80	0.750	0.254
2	**25.98**	**0.759**	**0.241**
3	25.76	0.756	0.247
4	25.44	0.747	0.258

Table 7. **Texture limit.** We train 40K primitives at different texture limits, K, on the Mip-NeRF360 scenes.

and TRIPOD, and to 1 million for TABLE. We use 1 million primitives for TABLE because using 3 million resulted in unstable behavior, and the PSNR dropped significantly.

MiniSplatting [12]. We ran MiniSplatting with their recommended indoor and outdoor settings for each respective scene. To obtain the two low memory comparisons shown in Figure 7, we set the sampling factor to 0.2 and 0.25, respectively.

NeST-Splatting [59]. The densification strategy in NeST-Splatting does not allow setting a maximum number of primitives. On our custom GRAFFITI scene, this resulted in an out-of-memory error even with a 48GB GPU. The results of NeST-Splatting on the custom dataset reported in the paper were obtained by doubling the interval between densification steps for the GRAFFITI scene alone. To attain the models shown in Figure 6 of the main paper, with exactly 40K primitives, we added a maximum number to their densification. We initialize from a COLMAP point cloud sampled using farthest point sampling [27] to 20K points, matching our initialization. The storage memory was computed using 16-bit floating point numbers for the hash-grid parameters.

10.2. Additional Results

We list the data points from Figure 5 of the main paper in Table 5. We include per-scene metrics for our method at
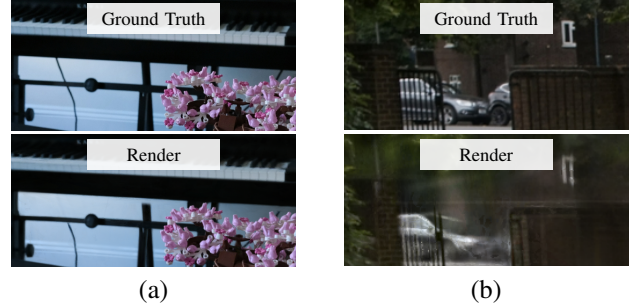

Figure 9. **Geometric inaccuracies.** (a) At a low number of primitives (40K) trained on the BONSAI scene, the nexels miss thin structures. (b) With 400K primitives trained on the BICYCLE scene, the background surfels optimize to an incorrect geometry and the textures learn to blur the details.

Routine	Indoor	Outdoor
Collection Pass	11.8 ms	8.1 ms
Texturing Pass	12.7 ms	8.4 ms
Total	24.5 ms	16.4 ms

Table 8. **Detailed timings.** We show timings for each pass averaged across the Mip-NeRF360 indoor and outdoor scenes with 400K primitives.

different settings in Table 6.

We also conduct an experiment to determine the best value of K in Table 7. The metrics sharply improve when adding in any texture at all (i.e. going from $K = 0$ to $K = 1$) and reach a peak at $K = 2$. Surprisingly, increasing K beyond 2 leads to degradation in quality. Empirically, we observe noise in the background of test renders at higher values of K, suggesting that primitives which were never textured at training views are textured at test views.

11. Further Discussion

11.1. Texturing Strategy

Limiting the number of textured primitives to K has a number of consequences which we outline. Firstly, the rendering speed becomes a compromise between that of rendering with per-primitive colours and that of texturing every primitive. The runtime of the texturing step is determined by the number of pixels and the number of per-pixel primitives that are textured, which is fixed.

In the forwards pass, the textured surfels determine the majority of the appearance. However, the non-textured surfels still make small improvements to the render in cases where the texture's appearance is not sufficient. This is useful around areas with strong reflections and transparent materials, as our rendering formulation does not encapsulate

4

full lighting effects.

In the backwards pass, allowing non-textured primitives to contribute to the render allows them to receive gradient flow during back-propagation. This is important for early training iterations where the gradient moves primitives towards the underlying scene geometry, regardless of whether the primitive was textured. The neural texture does not need to be rendered at multiple overlapping primitives to guide its optimization. Thanks to the volumetric nature of the neural field, its gradient flow is the same whether it is rendered by several coinciding surfels or by a single surfel with the same total weight.

11.2. Metrics

We primarily show LPIPS throughout the paper as we find that it best corresponds to the reconstruction quality of fine texture details, which is a weakness of Gaussian splatting. It also generally corresponds better to human preference, as that was its original purpose [58]. We note that the LPIPS formula used by virtually all Gaussian splatting works has a minor discrepancy from the one used by Zhang et al. [58]. Specifically, the official 3DGS repository scales its images to $[0, 1]$ before computing the LPIPS, while the LPIPS library used expects images scaled to $[-1, 1]$. Nonetheless, the scores returned using this appear to be reasonable for comparison purposes so long as all baselines are calculated in the same way. To be consistent with the rest of the field, we use the same LPIPS calculation as 3DGS.

11.3. Limitations

As mentioned in the main text, nexels are not able to run at interactive rates on mobile devices or low-end GPUs. We show timings of the individual rendering passes in Table 8. The neural texture step currently uses tensor core operations and would be much slower if implemented in WebGL or OpenGL.

Figure 8 includes qualitative examples of the limitations mentioned in the main paper, particularly unseen regions and motion blur. One more limitation is the geometric accuracy of the representation, particularly in background regions. If the geometry does not optimize or densify correctly, then the neural texture is of little help. We show two failure cases in Figure 9. In the first, the sparse geometries are not able to optimize towards thin structures. In the second, the background geometries optimize to a local minimum, resulting in blurry textures.