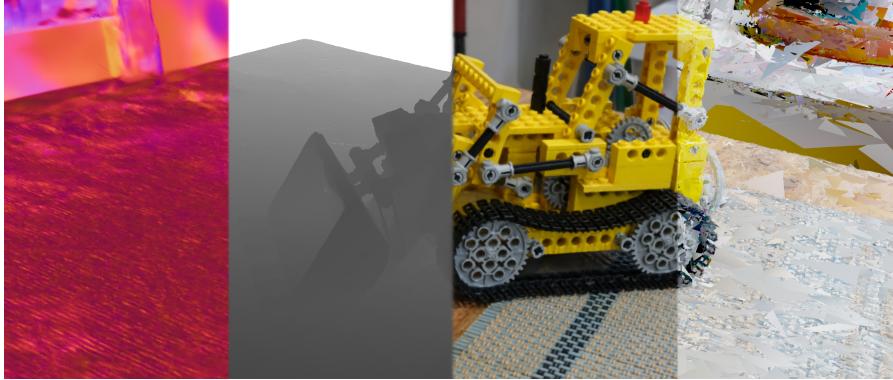

000 RADIANT TRIANGLE SOUP WITH SOFT CONNECTIVITY 001 FORCES FOR 3D RECONSTRUCTION AND NOVEL VIEW 002 SYNTHESIS 003 004

006 **Anonymous authors**

007 Paper under double-blind review



022 Figure 1: Optimizing Radiant Triangle Soup (RTS) produces high quality 3D models from captured images.
023 Above, from left to right, are the rendered normal map, rendered depth map, rendered image, and the direct
024 rasterization of the triangle primitives, as if they were opaque.

025 ABSTRACT

027 We introduce an inference-time scene optimization algorithm utilizing triangle
028 soup, a collection of disconnected translucent triangle primitives, as the repre-
029 sentation for the geometry and appearance of a scene. Unlike full-rank Gaussian
030 kernels, triangles are a natural, locally-flat proxy for surfaces that can be connected
031 to achieve highly complex geometry. When coupled with per-vertex Spherical
032 Harmonics (SH), triangles provide a rich visual representation without incurring
033 an expensive increase in primitives. We leverage our new representation to incor-
034 porate optimization objectives and enforce spatial regularization directly on the
035 underlying primitives. The main differentiator of our approach is the definition
036 and enforcement of soft connectivity forces between triangles during optimization,
037 encouraging explicit, but soft, surface continuity in 3D. Experiments on repres-
038 entative 3D reconstruction and novel view synthesis datasets show improvements
039 in geometric accuracy compared to current state-of-the-art algorithms without
040 sacrificing visual fidelity.

041 1 INTRODUCTION

044 Gaussian Splatting (GS) methods are effective at Novel View Synthesis (NVS). However 3D Gaussian
045 Splatting (3DGS) (Kerbl et al., 2023) fails to accurately model the geometry of scene surfaces. 3D
046 Gaussians are by design smooth, unbounded volumetric primitives, which are inherently ill-suited for
047 representing flat surfaces and sharp boundaries. Optimization for novel view synthesis involves alpha-
048 blending several overlapping Gaussians, none of which have to be located on the underlying surfaces.
049 As a result, optimizing a representation of 3D Gaussian kernels for image synthesis commonly causes
050 floaters or blurry artifacts in the scene.

051 Several authors (Huang et al., 2024; Guédon & Lepetit, 2024; Chen et al., 2024a; Dai et al., 2024)
052 have proposed flattening the kernels, with either strict 2D Gaussians (Huang et al., 2024) or flattened
053 3D Gaussians via loss regularization (Guédon & Lepetit, 2024; Chen et al., 2024a), to better model
thin surfaces. These methods have shown that modifying the underlying primitives and enforcing

optimization objectives on rendered geometry leads to improvements in geometric accuracy. While flat primitive help with surface alignment, diffuse ellipsoids still fail to accurately model depth discontinuities (see Fig. 8).

Triangles are the fundamental primitives in computer graphics because they are necessarily planar and convex. They can approximate any surface as piecewise planar at any tolerance level, can directly model sharp discontinuities, and can be rendered rapidly without approximations.

We present a new methodology for 3D reconstruction, which we named Radiant Triangle Soup (RTS) that enables gradient-based optimization of a Radiance Field (RF) using translucent triangle primitives as the scene representation (see Fig. 1). We provide a complete framework, including differentiable mechanisms for rasterization, as well as non-differentiable mechanisms for initialization, pruning, and densification. Our experiments show RTS achieves competitive results in terms of both appearance and geometry (see Fig. 7).

Conceptually, RTS provides a feature that is not supported by any other GS scene representation: an avenue for explicit information sharing among 3D among primitives. Conventional Gaussian kernels interact via alpha-blending when rendered onto common pixels. Back-propagating from the loss at each pixel to the primitives is the only means of coordination across Gaussians. Using multiple images from various viewpoints gives rise to several indirect constraints on the primitives, but direct constraints are currently absent.

Conversely, RTS is the first framework to enable direct information exchange among neighboring primitives. We formulate connectivity losses between neighboring triangle edges, allowing for a more direct and effective coordination of, and constraint on, primitive behavior throughout the optimization process. Even though using triangles as primitives alone achieves better geometric accuracy compared to Gaussian surfels, encouraging connectivity during optimization leads to more accurate surfaces with less floating artifacts (see Table 3).

Color expressivity on a per-primitive basis is another dimension where the RTS representation excels. It is common practice to use a single color parametrization for each primitive. However, this may require large numbers of overlapping primitives in order to reconstruct details in the surfaces of the scene. Using triangles as the representation allows for each vertex to encode a separate color (in the form of Spherical Harmonics), leading to more expressive primitives through bilinear interpolation.

To summarize, the work presented here:

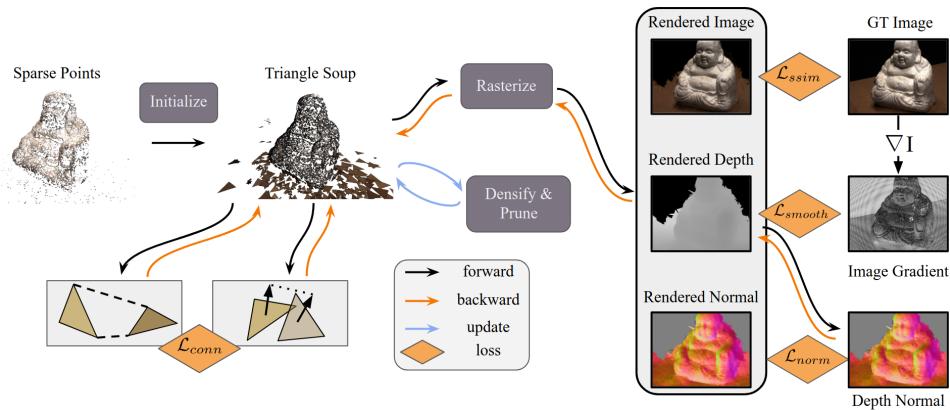
- Develops a new scene representation through alpha-blending of triangle primitives.
- Introduces explicit 3D forces between primitives to encourage soft connectivity.
- Increases primitive expressivity via multi-color encoding.

2 RELATED WORK

We begin this section with Gaussian splatting formulations that favor surfaces and continue with methods that rely on non-Gaussian primitives. Surveys on other aspects of Gaussian splatting, omitted due to space limitations, include (Bao et al., 2025; Dalal et al., 2024; Luo et al., 2024; Wu et al., 2024). The seminal work of Kerbl et al. (2023) on 3D Gaussian Splatting introduced an explicit alternative to NeRF (Mildenhall et al., 2020) that is able to achieve high-quality rendering at much higher speed. 3DGS relies on interleaved differentiable optimization and non-differentiable density control of the explicit representation. The optimization process decreases view synthesis errors for one of the training images at each iteration and density control guides the placement of primitives.

2D Gaussian Splatting (2DGS) (Huang et al., 2024) modified 3DGS to prioritize the reconstruction of surfaces, rather than volumetric material. This was accomplished by collapsing the 3D GS to 2D discs by setting the minimum eigenvalue of the Gaussian to 0. The authors of SuGaR (Guédon & Lepetit, 2024) and PGSR (Chen et al., 2024a) devised similar techniques for aligning 2D Gaussian with the surfaces via regularization and a multi-view loss, respectively. Gaussian Surfels (Dai et al., 2024) rely on monocular surface normal estimates and a normal-depth consistency loss. Gaussian Opacity Fields (GOF) (Yu et al., 2024) extract surfaces as the zero-level set of 3D Gaussians. 3D-Half-Gaussian Splatting (3D-HGS) (Li et al., 2024) enables the representation of perfectly planar surfaces and hard

108
109
110
111
112
113
114
115
116
117
118
119
120
121



122 Figure 2: We begin with a set of sparse points. From these points we initialize a set of triangles
123 and compute nearest neighbor connections for each triangle edge. During optimization, we render
124 an image, depth map, and normal map for each view. We compute both 2D loss (over all output
125 renderings), as well as 3D loss (directly over the primitive connections). Adaptive densification
126 is non-differentiable and is performed at set intervals throughout the optimization. Triangle edge
127 neighbors are recomputed following adaptive densification.

128 edges by attaching a splitting plane to each Gaussian, aligning it to the local surface, and setting the
129 density of one of the half-spaces to 0.

130
131 Representations based on NeRF and 3DGS have found great success, but have also inspired re-
132 searchers to seek alternatives. NeuRBF (Chen et al., 2023) utilizes Radial Basis Functions (RBFs) to
133 overcome limitations of NeRF due to the global nature of its MLP and features. GES (Hamdi et al.,
134 2024) is based on an explicit representation which replaces the Gaussian kernel with a Generalized
135 Exponential Function, overcoming the low-pass effect of the Gaussian and thus requiring fewer
136 primitives to represent the scene. Similar approaches based on smooth, non-Gaussian kernels include
137 DARB-Splatting (Arunan et al., 2025), SolidGS (Shen et al., 2024), Beyond Gaussians (Chen et al.,
138 2024b) and Deformable Beta Splatting (Liu et al., 2025).

139 Besides the methods that force their Gaussians to be planar (Huang et al., 2024; Guédon & Lepetit,
140 2024; Chen et al., 2024a), but still diffuse, there are others that represent the scene with collections
141 of planar primitives. Zanjani et al. (2025) initially use Gaussian splats to model the scene and then
142 merge them into 3D planes, which are abundant in indoor scenes. PlanarSplatting (Tan et al., 2025)
143 is also designed for indoor scenes using planes, initialized via monocular depth estimation, as the
144 only primitives. TRIPS (Franke et al., 2024) is based on the principle that point primitives can be
145 rasterized into an image pyramid from which the appropriate layer can be selected according to the
146 size of the projected point. Holes can be filled by a small network yielding accurate, crisp renderings.
147 Triangle Splatting (Held et al., 2025a) uses triangles with diffuse edges as primitives, but does not
support any mechanism for them to interact directly with each other.

148 Non-planar primitives were introduced by BG-Triangle (Wu et al., 2025) which uses Bézier Gaussian
149 triangles that are effective near boundaries, but comes at the cost of operating on complex, non-
150 planar primitives that are hard to render. Quadratic Gaussian Splatting (Zhang et al., 2025) uses
151 deformable quadratic surfaces as primitives and geodesic, instead of Euclidean, distance-based
152 density distributions that adapt to the curvature of the primitives.

153 Another class of methods rely on volumetric primitives. LinPrim (von Lützow & Nießner, 2025) is
154 based on linear solid primitives with triangular faces and performs volumetric rendering by computing
155 the entrance and exit points of each ray through the primitives. 3D Convex Splatting (3DCS) (Held
156 et al., 2025b) was inspired by the limitation of GS that requires very large numbers of primitives to
157 model hard edges and flat surfaces due to the diffuseness of the Gaussians. As the name suggests, the
158 scene is represented by a set of polyhedral convexes which undergo volumetric rendering, pruning
159 and splitting operations. The concepts of smoothness and sharpness introduced by 3DCS have
160 inspired our diffuseness (see Section 3). Radiant Foam (Govindarajan et al., 2025) enables modeling
161 light transport phenomena, like reflection and refraction, by tessellating the space into Voronoi
cells and iteratively optimizing the positions of the Voronoi vertices. All methods in this paragraph

162 rely on different forms of volumetric elements with planar faces. **Computing the intersections of**
 163 **each primitive with a ray requires multiple ray-triangle intersections, compared to the single**
 164 **ray-triangle intersection required by RTS.**

165 A recent trend in the literature has been joint optimization of two representations: one for synthesizing
 166 novel views and one that is more faithful to the surfaces (Choi et al., 2024; Jiang et al., 2025). MILO
 167 (Guédon et al., 2025) maintains Gaussians that are alpha-blended for view synthesis and a watertight
 168 mesh without texture that captures the geometry of the scene.

169 A few common themes emerge by analyzing the above methods. Most advocate the use of bounded
 170 primitives to enhance their ability to represent sharp edges and flat surfaces with small numbers of
 171 primitives. Our primitives are bounded but have diffuse boundaries to facilitate optimization. No
 172 other method, however, has a mechanism for direct inter-primitive communication like RTS.

174 3 METHOD

175 Given a set of images together with camera poses and a
 176 set of sparse points \mathcal{S} computed via Structure-From-Motion
 177 (Schönberger & Frahm, 2016), we construct an explicit scene
 178 representation using triangle primitives, the parameterization
 179 and initialization of which we discuss in Sections 3.1 and 3.2,
 180 respectively. The triangles are endowed with diffuse bound-
 181aries, similar to 3DCS (Held et al., 2025b) (Section 3.3). To
 182 model the surfaces in the scene, we render into each camera
 183 the image, depth, and normal map of the triangles through
 184 alpha-blending (see Section 3.4). The triangle parameters
 185 are directly updated via back-propagation after computing
 186 losses between the rendering of the scene and the ground
 187 truth image from each view. We also include 2D losses on
 188 the rendered depth and normal maps (see Section 4). In our
 189 representation, primitives maintain *soft* connectivity with
 190 their neighbors, discussed in Section 3.5, for which we com-
 191 pute additional loss directly over the connection orientations (see Section 4). Similar to previous
 192 works (Kerbl et al., 2023; Huang et al., 2024; Guédon & Lepetit,
 193 2024; Chen et al., 2024a; Yu et al., 2024), we develop a strategy for adaptive density control in order to facilitate the addition and removal
 194 of primitives in the scene (see Section 3.6). We show an overview of our algorithm in Fig. 2.

196 3.1 PARAMETERIZATION

197 From the set \mathcal{S} of sparse points, we first create an initial set of triangles \mathcal{T} . The triangle primitives
 198 $t_n \in \mathcal{T}$ in our scene representation are parametrized with,

$$201 \quad t_n = \{\mu, \Delta, s, R, \alpha, \sigma\} \quad (1)$$

202 where μ is the incenter of the triangle, $\Delta = [\delta^0 \quad \delta^1 \quad \delta^2]$ are the per-vertex Spherical Harmonics,
 203 $s = [s^0 \quad s^1 \quad s^2]$ are the scales, R is the 3×3 rotation matrix, α is the opacity, and σ is the diffuse
 204 scalar, discussed in Section 3.3. **This parameterization uniquely defines each triangle primitive.**

207 3.2 INITIALIZATION

209 The coordinates of each sparse SfM point are used as the initialization for the incenter, and the point
 210 colors as the initialization for the zero-component of the Spherical Harmonics for all three vertices,
 211 which are then optimized separately. Similar to previous works (Huang et al., 2024; Kerbl et al.,
 212 2023), the scales for our primitives are computed based on the average distance to the three nearest
 213 neighboring points. Each triangle starts as equilateral, using the scales to parameterize the distance
 214 of each vertex from the incenter of the triangle μ along the bisectors of the angles (see Fig. 3). The
 215 diffuse scalar σ is also set as a function of the average distance to the three nearest neighboring points
 (Please see Eq. S.2). **The rotation of each triangle is initialized with a random rotation matrix.**

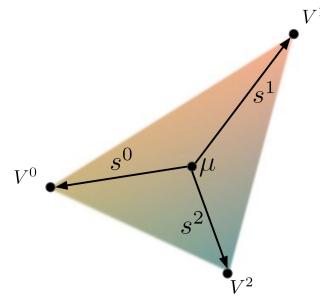
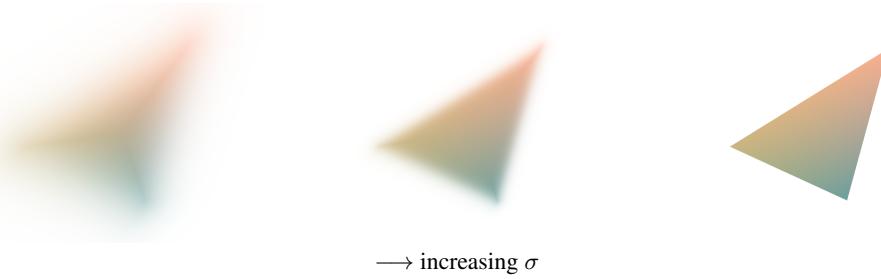


Figure 3: Triangle parameterization. Each triangle is parameterized by the incenter μ , and three scales $[s^0 \quad s^1 \quad s^2]$. These parameters, along with the rotation matrix R , define the coordinates of each vertex V^j .

216
217
218
219
220
221
222
223
224



225
226 Figure 4: Controlling triangle diffuseness with σ . Small values of σ (**left**) result in diffuse triangles
227 with extended range of influence. Large values of σ (**right**) result in sharp triangles with no blurring
228 across edges.

229
230

3.3 DIFFUSE PRIMITIVE BOUNDARIES

231 Alpha-blending is the primary driving force of optimization for all splatting frameworks. Blending
232 across primitives facilitates optimization via gradient descent and encourages, among other behaviors,
233 movement in primitive position and orientation. It is challenging in practice to optimize a scene
234 with fully opaque primitives that do not smoothly blend with one another. It is therefore desirable
235 to make the triangle primitives diffuse near the edges. Taking inspiration from 3DCS (Held et al.,
236 2025b), we parameterize the response from each triangle as a function of the signed distance from
237 the ray-triangle intersection to the nearest edge,

238
239
$$w_\sigma = \frac{1}{1 + e^{(\sigma l)}} \quad (2)$$

240 where w_σ is the diffuse weight, l is the signed distance between the intersection and the nearest
241 triangle edge in the plane of the triangle, and the scalar σ is an optimizable parameter of each triangle
242 that controls the level of diffuseness. As σ increases, the boundary of the triangle becomes less
243 diffuse, creating sharper primitive renderings (see Fig. 4). For the signed distance, $l < 0$ occurs when
244 the intersection point lies outside of the triangle boundary. This formulation slightly deviates from
245 that of 3DCS, as the signed distance is computed directly, compared to their smooth approximation.

246
247

3.4 RASTERIZATION

248 Throughout this paper, we use barycentric coordinates, $\lambda = [\lambda^0, \lambda^1, \lambda^2]^T$, to rasterize triangles.
249 Please see Section S.1 for more details. With the barycentric coordinates of the ray-triangle intersec-
250 tions and the diffuse weight computed, we interpolate the color for the current pixel w.r.t. the colors
251 of each vertex weighted by the barycentric coordinates, $c_n = [c^0 \ c^1 \ c^2] \lambda$, where c^0 , c^1 , and
252 $c^2 \in \mathbb{R}^3$ are computed from the SH components of each vertex, and λ are the barycentric coordinates
253 of the intersection point.

254 We aggregate the contribution to the current output pixel, i , from each intersected primitive, $c_i =$
255 $\sum w_n c_n$, where $w_n = \alpha w_\sigma T$, and $T = \prod_{j=1}^{i-1} (1 - \alpha_j)$ is the transmittance for the current primitive
256 (Kerbl et al., 2023).

257 Previous works (Huang et al., 2024; Chen et al., 2024a; Yu et al., 2024) provide two methods for
258 rendering per-pixel depth; (1) computing the average weighted intersection depth of all traversed
259 primitives (mean depth), and (2) using only the depth of the primitive that causes the transmittance T
260 to exceed 0.5 (median depth). In our work, we use median depth. Using the mean depth encourages
261 the formation of many translucent layers of primitives. Rendering median depth removes this blending
262 and helps guide the formation of surfaces. We directly use d from Eq. S.2 for the depth of each
263 intersection.

264 Additionally, we compute the surface normal, \hat{n}_n , for each triangle as the cross-product between
265 two edges. Unlike the rendered depth, to render per-pixel surface normals, we follow previous work
266 (Huang et al., 2024; Chen et al., 2024a) and alpha-blend all the surface normals of all intersected
267 primitives. The intuition is that through a normal consistency loss (see Section 4), the blended
268 normals must align with the normals computed from the median depth map. This encourages all

270 the intersected triangles to align (and eventually collapse onto) the median surface. Thus, we render
 271 per-pixel surface normals through alpha-blending, $\vec{n}_i = \sum w_n \hat{n}_n$.
 272

273 3.5 CONNECTIVITY 274

275 To enable soft connectivity among primitives and encourage
 276 flat, connected surfaces, we add a connectivity term to the
 277 optimization objective. For each triangle edge, we assign, and
 278 periodically update, a connection to the nearest neighboring
 279 triangle edge. Taking inspiration from the energy functions
 280 introduced in 3D scene flow estimation (Vogel et al., 2015),
 281 the connectivity term of the loss increases according to the
 282 distance between their vertices and the inner product between
 283 their normals, discussed in Section 4.

284 Naively connecting with the nearest edge without considering
 285 the relative orientations may cause connections that would
 286 require large changes in rotation to either triangle, leading to
 287 undesirable behaviors during optimization. To prevent this behavior, connections are only established
 288 if the triangle edges are "facing" each other. In practice, we use the inner product between the unit
 289 vectors orthogonal to the triangle edges (in the plane of each triangle) as the criterion for establishing
 290 connections. We provide an example in Fig. 5 where a connection to the highlighted edge of triangle
 291 D is valid, while the other two connections would cause large rotations.
 292

293 3.6 ADAPTIVE DENSITY CONTROL 294

295 We perform adaptive density control through the process of cloning, splitting and pruning triangles.
 296 Previous works (Kerbl et al., 2023; Huang et al., 2024; Guédon & Lepetit, 2024; Chen et al., 2024a)
 297 perform the cloning and splitting procedures by duplicating primitives conditioned on scale and
 298 position gradients. In order to properly split large triangles, we must split them into four sub-
 299 triangles. During early iterations in the optimization, when triangles are split, the sub-triangles
 300 move independently to better align with surfaces in the scene. During the later stages when most
 301 of the surfaces have formed, the triangles split to enable more detailed surface color representation,
 302 remaining relatively attached via the connectivity forces. See Fig. 6 for a visualization of the splitting
 303 procedure.

304 Similar to previous work, we directly clone small primitives selected
 305 for densification. Much like 2DGS (Huang et al., 2024), our method
 306 does not directly rely on the gradient of the projected 2D primitive
 307 center. Instead of computing an approximation via projecting 3D
 308 gradients into the camera plane (Huang et al., 2024), we directly
 309 condition densification on the magnitude of the incenter gradient
 $\nabla \mu$ in 3D.

310 To remove uninformative triangles, we prune primitives that meet
 311 the following criteria: (i) triangles that are transparent ($\alpha < 0.05$), (ii) triangles with one or fewer
 312 edge connections, (iii) triangles that do not intersect the camera frustum of at least three views after
 313 each epoch.

314 4 OPTIMIZATION 315

316 Our objective function comprises terms computed in 2D (on the images plane), and in 3D (on the
 317 triangles).

318 4.1 RENDERING LOSSES 319

320 Following previous work (Kerbl et al., 2023; Huang et al., 2024), we compute the SSIM between the
 321 rendered and input images, \mathcal{L}_{ssim} , and apply a normal consistency loss (Huang et al., 2024), \mathcal{L}_{norm} ,
 322

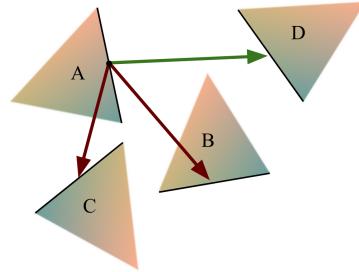


Figure 5: Examples of edge association for connecting nearby triangles. For simplicity, it is assumed the triangles are *co-planar* in this figure.

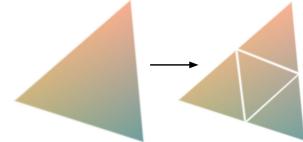
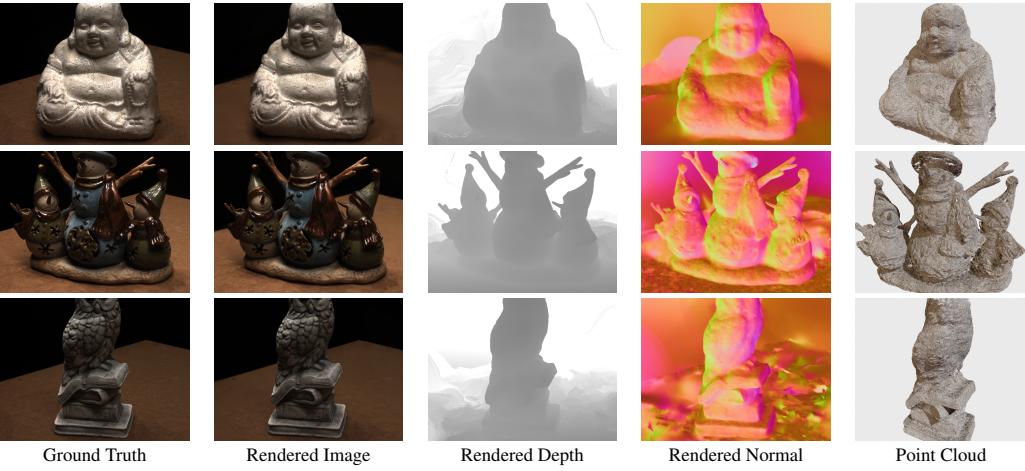


Figure 6: Triangles are *split* with interpolated colors per-vertex.



324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377

Figure 7: Qualitative results on the DTU dataset (Aanæs et al., 2016). RTS estimates high-quality geometry, maintaining thin structures, such as the carrot noses and branches (middle row) and the book edges (bottom row), while effectively modeling smooth surfaces, such as the Buddha (top row).

to help locally align the triangles with the rendered surface. Please see Section S.3 the supplement for further details.

From the unsupervised depth estimation literature (Chang et al., 2022; Godard et al., 2017; Mahjourian et al., 2018), we adopt a smoothness term on the rendered depth map conditioned on the gradient of the input image. This penalizes large gradients in the rendered depth map where we have small gradients in the input image,

$$\mathcal{L}_{smooth} = \frac{1}{N} \sum_{i,j} \|\partial_x D_{i,j}\| e^{\|\partial_x I_{i,j}\|} + \|\partial_y D_{i,j}\| e^{\|\partial_y I_{i,j}\|} \quad (3)$$

4.2 SCENE LOSS

To encourage connectivity between primitives, we penalize the mean of the L_2 distance between the connected vertex pairs of neighboring triangle edges,

$$\mathcal{L}_{conn} = \sum_{a \in \Omega} \frac{1}{2} (\|V_a^1 - V_b^1\|_2 + \|V_a^2 - V_b^2\|_2) + (1 - \hat{n}_a^T \hat{n}_b) \quad (4)$$

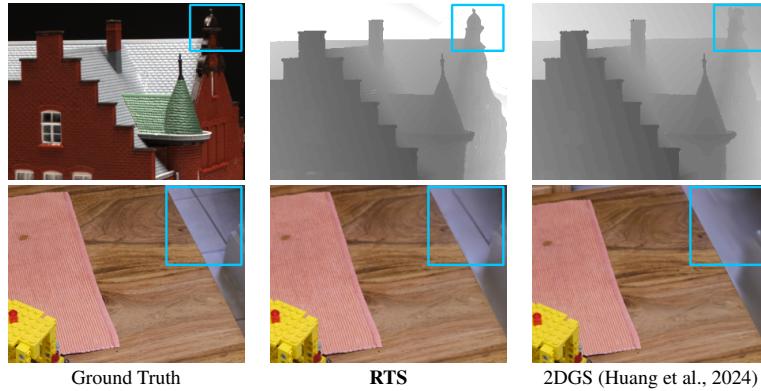
where V_a^j and V_b^j are the j^{th} vertex pair of connected edges a and b , respectively. \hat{n}_a and \hat{n}_b are the normals of the connecting triangles and Ω is the set of all triangles that intersect the current camera frustum. **The normal regularization is a soft penalty that encourages connected triangles to have similar normals to better align with the surfaces in the scene.** Applying this loss to invisible triangles without rendering losses leads to over-smoothing.

The final objective is a weighted summation of all terms:

$$\mathcal{L} = \omega_0 \mathcal{L}_{ssim} + \omega_1 \mathcal{L}_{norm} + \omega_2 \mathcal{L}_{smooth} + \omega_3 \mathcal{L}_{conn} \quad (5)$$

Table 1: Chamfer distance evaluation on scenes from DTU (Aanæs et al., 2016). Following previous literature, we average the accuracy and completeness Chamfer distances on the widely used evaluation set. Chamfer distances are measured in millimeters. The best results are in boldface and the second best are underlined.

Method	DTU															
	24	37	40	55	63	65	69	83	97	105	106	110	114	118	122	Mean (mm)↓
3DGs (Kerbl et al., 2023)	2.14	1.53	2.08	1.68	3.49	2.21	1.43	2.07	2.22	1.75	1.79	2.55	1.53	1.52	1.50	1.96
SuGaR (Guédon & Lepetit, 2024)	1.47	1.33	1.13	0.61	2.25	1.71	1.15	1.63	1.62	1.07	0.79	2.45	0.98	0.88	0.79	1.33
2DGs (Huang et al., 2024)	0.48	0.91	0.39	0.39	1.01	0.83	0.81	1.36	1.27	0.76	0.70	1.40	0.40	0.76	0.52	0.80
GOF (Yu et al., 2024)	0.50	0.82	0.37	0.37	1.12	0.74	0.73	1.18	1.29	0.68	0.77	0.90	0.42	0.66	0.49	0.74
Gaussian Surfs (Dai et al., 2024)	0.66	0.93	0.54	0.41	1.06	1.14	0.85	1.29	1.53	0.79	0.82	1.58	0.45	0.66	0.53	0.88
PGSR (Chen et al., 2024a)	0.36	0.57	0.38	0.33	0.78	0.58	0.50	1.08	0.63	0.59	0.46	0.54	0.30	0.38	0.34	0.52
TriangleSplatting (Held et al., 2025a)	0.98	1.07	1.07	0.51	1.67	1.44	1.17	1.32	1.75	0.98	0.96	1.11	0.56	0.93	0.72	1.06
RTS	<u>0.42</u>	<u>0.61</u>	0.74	0.39	<u>0.53</u>	0.86	<u>0.70</u>	0.84	<u>0.72</u>	0.37	<u>0.68</u>	<u>0.87</u>	0.34	<u>0.58</u>	<u>0.44</u>	<u>0.61</u>



378
 379
 380
 381
 382
 383
 384
 385
 386
 387
 388
 389
 390
 391 Figure 8: Qualitative comparison between RTS and 2DGS (Huang et al., 2024) on; **Top** - *scan024*
 392 from the DTU dataset (Aanæs et al., 2016), **Bottom** - *kitchen* from the mip-NeRF 360 dataset (Barron
 393 et al., 2022). RTS is substantially more precise at estimating the geometry at discontinuities and
 394 rendering fine details, shown in the areas marked by the blue rectangles.

5 EXPERIMENTS

5.1 IMPLEMENTATION DETAILS

400 We implement the majority of our RTS framework in Python using PyTorch (Paszke et al., 2019).
 401 For rasterization, we develop custom CUDA kernels for both the forward and backward pass.
 402 We run all our experiments on a single NVIDIA RTX A6000. For our loss weights, we choose
 403 $\omega = [1.0 \ 0.05 \ 50.0 \ 1000.0]$ empirically.

404 Following state-of-the-art Multi-View Stereo methods (Yao et al., 2018; Yang et al., 2022; Mi et al.,
 405 2022), we directly generate a 3D point cloud for geometric evaluation from the rendered depth maps
 406 *without performing any TSDF fusion*. To generate each point cloud, we use simple heuristic filtering
 407 on each depth map, similar to the post-processing presented in GBiNet (Mi et al., 2022). For each
 408 depth map, we measure the reprojection error of the depth values at every pixel using neighboring
 409 views and filter pixels based on this error. All depth estimates with a low reprojection error are
 410 back-projected to 3D points, forming the combined point cloud.

5.2 EVALUATION

414 We test our framework on the DTU dataset (Aanæs et al., 2016), an indoor dataset that contains
 415 images of 124 scenes taken from a camera mounted on an industrial robot arm. All scenes share the
 416 same camera trajectories, with ground-truth point clouds captured via structured light.

417 We evaluate our new approach on the DTU dataset and record the Chamfer distance in Table 1.
 418 We show competitive results alongside the leading state-of-the-art planar GS methods. Across the
 419 test set, RTS is the most geometrically accurate in several scenes and second most overall. While

422 Table 2: Novel View Synthesis on all scenes from the mip-NeRF 360 dataset (Barron et al., 2022).

Method	Outdoor			Indoor		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
NeRF (Mildenhall et al., 2020)	21.46	0.458	0.515	26.84	0.790	0.370
Deep Blending (Hedman et al., 2018)	21.54	0.524	0.364	26.40	0.844	0.261
Instant NGP (Müller et al., 2022)	22.90	0.566	0.371	29.15	0.880	0.216
MipNeRF360 (Barron et al., 2022)	24.47	0.691	0.283	31.72	0.917	0.180
SuGaR (Guédon & Lepetit, 2024)	22.93	0.629	0.356	29.43	0.906	0.225
3DGS (Kerbl et al., 2023)	24.64	0.731	0.234	30.41	0.920	0.189
2DGS (Huang et al., 2024)	24.34	0.717	0.246	30.40	0.916	0.195
GOF (Yu et al., 2024)	24.82	0.750	0.202	30.79	0.924	0.184
PGSR (Chen et al., 2024a)	24.76	0.752	0.203	30.36	0.934	0.147
3DCS (Held et al., 2025b)	24.07	0.700	0.238	31.33	0.927	0.166
TriangleSplatting (Held et al., 2025a)	24.27	0.722	0.217	30.80	0.928	0.160
RTS	21.41	0.657	0.349	30.28	0.921	0.130

432
433 Table 3: Ablation study on the contribution of the 3D loss terms using the entire DTU evaluation set.
434 Here, λ_c is the loss weight for the connectivity term, L_{conn} , and λ_s is the loss weight for the depth
435 smoothness term L_{smooth} . The best results are in boldface and the worst are underlined.

Method	Acc.(mm) \downarrow	Comp.(mm) \downarrow	SSIM \uparrow	PSNR \uparrow	LPIPS \downarrow	Primitives (K) \downarrow
w/o \mathcal{L}_{smooth} & \mathcal{L}_{conn}	0.66	0.68	0.912	30.58	0.227	218
w/o \mathcal{L}_{smooth}	0.65	<u>0.71</u>	0.921	32.13	0.212	244
w/o \mathcal{L}_{conn}	<u>0.67</u>	0.69	<u>0.908</u>	<u>30.15</u>	0.231	224
full ($\lambda_c = 10.0$, $\lambda_s = 0.8$)	0.64	0.70	0.918	31.68	0.213	249
full ($\lambda_c = 300.0$, $\lambda_s = 20.0$)	0.61	0.63	0.910	30.87	<u>0.232</u>	244
full ($\lambda_c = 1000.0$, $\lambda_s = 50.0$)	0.59	0.62	0.909	30.57	<u>0.232</u>	<u>297</u>

441
442 PGSR demonstrates impressive reconstruction results, the algorithm utilizes a full suite of multi-view
443 objective functions that significantly improve the geometric reconstruction quality. We provide
444 qualitative results on the DTU dataset in Fig. 7, showing visualizations of the rendered images, depth
445 maps, normal maps, and final point clouds for three scenes. In Fig. 8 (top), we provide a comparison
446 of depth map renderings between RTS and 2DGS (Huang et al., 2024). RTS is able to reconstruct
447 fine details on the surfaces of objects that are typically blurred with Gaussian representations.

448 We show additional results on the mip-NeRF 360 dataset (Barron et al., 2022). Following the protocol
449 specified by Barron et al. (2022), we separate the images in each scene, taking every eighth image as a
450 test image and training on the remaining. As standard evaluation, we report PSNR, SSIM, and LPIPS
451 (Zhang et al., 2018) metrics. We show results on all indoor and outdoor scenes in Table 2, as well as
452 per-scene results in Table S.2. For indoor scenes, RTS shows results on par with the state-of-the-art
453 methods, having the leading LPIPS score, the fourth overall SSIM score, and a highly competitive
454 PSNR score among all listed methods. On two outdoor scenes (treehill and flowers), RTS is limited
455 in reconstructing extremely distant background foliage, impacting the overall metrics (please see
456 Fig. S.5 for qualitative results on these scenes). In Fig. 8 (bottom), we provide a comparison of novel
457 view synthesis between RTS and 2DGS. RTS is able to render fine textures in low visibility regions
458 in scenes as opposed to blurring with Gaussian representations.

459
460

5.3 ABLATIONS

461 We show an ablation evaluating the contributions of proposed loss terms in Table 3. Removing the
462 soft connectivity leads to a decrease in overall accuracy of the output models, while removing the
463 depth supervision negatively affects the completeness. The two supervision signals complement
464 each other, and we show that increasing the loss weights allows for tuning the framework for either
465 better novel view synthesis or geometry. Additionally, the results corresponding to the changes in
466 magnitude of the loss weights demonstrates the stability of RTS to changes in hyper-parameters.

467 To evaluate geometry, all previous planar GS algorithms utilize a GT foreground-background seg-
468 mentation mask when generating the final models. Using this mask to generate the final models for
469 evaluation removes the effects of floaters and inaccurate estimation near the surfaces being evaluated.
470 To portray a more grounded evaluation of geometry, we provide an ablation study in Table S.3 in
471 which we compute Chamfer distance of models without the use of any GT masks.

472
473

6 LIMITATIONS & CONCLUSIONS

474 In this work, we introduce a new scene representation, namely Radianc Triangle Soup (RTS). To the
475 best of our knowledge, we are the first to introduce explicit 3D forces between primitives in a splatting
476 framework, helping to coordinate the positioning of primitives to directly form surfaces. Modifying
477 the weights of these forces allows for tuning between 3D reconstruction quality and novel-view
478 synthesis quality. The main limitation of our current algorithm is its inability to extract watertight
479 meshes. Furthermore, due to the periodic nearest-neighbors search, there is a minor increase in
480 run-time proportional to the number of primitives. Please see Section S.2 for more details.

481 The introduction of Triangle Soup as the underlying representation for Radiance Fields is amenable
482 to future work in surface optimization. We plan to extend RTS with modified primitive connectivity
483 strategies and perform optimization over watertight meshes.

486
487

7 REPRODUCIBILITY

488 In order to ensure reproducibility, we supplement the description of our method in Section 3 with
489 broad implementation details in Section 5.1, as well as a complete account of hyper-parameter values
490 used in our experiments in Section S.2 of the Supplemental Material. We will make our code publicly
491 available to the research community, if the paper is accepted.

492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539

540 REFERENCES
541

542 Henrik Aanæs, Rasmus Ramsbøl Jensen, George Vogiatzis, Engin Tola, and Anders BJORHOLM Dahl.
543 Large-Scale Data for Multiple-View Stereopsis. *IJCV*, 2016.

544 Vishagar Arunan, Saeedha Nazar, Hashiru Pramuditha, Vinasirajan Viruthshaan, Sameera Ramas-
545 inghe, Simon Lucey, and Ranga Rodrigo. DARB-Splatting: Generalizing Splatting with Decaying
546 Anisotropic Radial Basis Functions. *arXiv preprint arXiv:2501.12369*, 2025.

547 Yanqi Bao, Tianyu Ding, Jing Huo, Yaoli Liu, Yuxin Li, Wenbin Li, Yang Gao, and Jiebo Luo. 3D
548 Gaussian Splatting: Survey, Technologies, Challenges, and Opportunities. *IEEE Transactions on*
549 *Circuits and Systems for Video Technology*, 2025.

551 Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-NeRF
552 360: Unbounded Anti-Aliased Neural Radiance Fields. In *CVPR*, 2022.

554 Di Chang, Aljaž Božič, Tong Zhang, Qingsong Yan, Yingcong Chen, Sabine Süsstrunk, and Matthias
555 Nießner. RC-MVSNet: Unsupervised Multi-View Stereo with Neural Rendering. In *ECCV*, 2022.

556 Danpeng Chen, Hai Li, Weicai Ye, Yifan Wang, Weijian Xie, Shangjin Zhai, Nan Wang, Haomin
557 Liu, Hujun Bao, and Guofeng Zhang. PGSR: Planar-based Gaussian Splatting for Efficient and
558 High-Fidelity Surface Reconstruction. *IEEE TVCG*, 2024a.

559 560 Haodong Chen, Runnan Chen, Qiang Qu, Zhaoqing Wang, Tongliang Liu, Xiaoming Chen, and
561 Yuk Ying Chung. Beyond Gaussians: Fast and High-Fidelity 3D Splatting with Linear Kernels.
562 *ArXiv*, 2024b.

563 564 Zhang Chen, Zhong Li, Liangchen Song, Lele Chen, Jingyi Yu, Junsong Yuan, and Yi Xu. NeuRBF:
565 A Neural Fields Representation with Adaptive Radial Basis Functions. In *ICCV*, 2023.

566 567 Jaehoon Choi, Yonghan Lee, Hyungtae Lee, Heesung Kwon, and Dinesh Manocha. MeshGS:
568 Adaptive Mesh-aligned Gaussian Splatting for High-quality Rendering. In *ACCV*, pp. 3310–3326,
569 2024.

570 571 Pinxuan Dai, Jiamin Xu, Wenxiang Xie, Xinguo Liu, Huamin Wang, and Weiwei Xu. High-quality
572 Surface Reconstruction using Gaussian Surfels. In *SIGGRAPH*, 2024.

573 574 Anurag Dalal, Daniel Hagen, Kjell G Robbersmyr, and Kristian Muri Knausgård. Gaussian Splatting:
575 3d Reconstruction and Novel View Synthesis, a Review. *IEEE Access*, 2024.

576 577 Linus Franke, Darius Rückert, Laura Fink, and Marc Stamminger. TRIPS: Trilinear Point Splatting
578 for Real-Time Radiance Field Rendering. *Computer Graphics Forum*, 43(2), 2024.

579 580 Clément Godard, Oisin Mac Aodha, and Gabriel J Brostow. Unsupervised Monocular Depth
581 Estimation with Left-Right Consistency. In *CVPR*, 2017.

582 583 Shrisudhan Govindarajan, Daniel Rebain, Kwang Moo Yi, and Andrea Tagliasacchi. Radiant Foam:
584 Real-Time Differentiable Ray Tracing. *ArXiv*, 2025.

585 586 Antoine Guédon and Vincent Lepetit. SuGaR: Surface-Aligned Gaussian Splatting for Efficient 3D
587 Mesh Reconstruction and High-Quality Mesh Rendering. In *CVPR*, 2024.

588 589 Antoine Guédon, Diego Gomez, Nissim Maruani, Bingchen Gong, George Drettakis, and Maks
590 Ovsjanikov. MILO: Mesh-In-the-Loop Gaussian Splatting for Detailed and Efficient Surface
591 Reconstruction. In *SIGGRAPH Asia*, 2025.

592 593 Abdullah Hamdi, Luke Melas-Kyriazi, Jinjie Mai, Guocheng Qian, Ruoshi Liu, Carl Vondrick,
594 Bernard Ghanem, and Andrea Vedaldi. GES : Generalized Exponential Splatting for Efficient
595 Radiance Field Rendering. In *CVPR*, pp. 19812–19822, 2024.

596 597 Peter Hedman, Julien Philip, True Price, Jan-Michael Frahm, George Drettakis, and Gabriel Brostow.
598 Deep Blending for Free-Viewpoint Image-Based Rendering. *ACM TOG*, 2018.

594 Jan Held, Renaud Vandeghen, Adrien Deliege, Abdullah Hamdi, Silvio Giancola, Anthony Cioppa,
595 Andrea Vedaldi, Bernard Ghanem, Andrea Tagliasacchi, and Marc Van Droogenbroeck. Triangle
596 Splatting for Real-Time Radiance Field Rendering. *ArXiv*, 2025a.
597

598 Jan Held, Renaud Vandeghen, Abdullah Hamdi, Adrien Deliege, Anthony Cioppa, Silvio Giancola,
599 Andrea Vedaldi, Bernard Ghanem, and Marc Van Droogenbroeck. 3D Convex Splatting: Radiance
600 Field Rendering with 3D Smooth Convexes. In *CVPR*, 2025b.

601 Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2D Gaussian Splatting
602 for Geometrically Accurate Radiance Fields. In *SIGGRAPH*, 2024.

603 Changjian Jiang, Kerui Ren, Lining Xu, Jiong Chen, Jiangmiao Pang, Yu Zhang, Bo Dai, and Mulin
604 Yu. HaloGS: Loose Coupling of Compact Geometry and Gaussian Splats for 3D Scenes. *arXiv*
605 preprint *arXiv:2505.20267*, 2025.

606

607 Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3D Gaussian
608 Splatting for Real-Time Radiance Field Rendering. *ACM TOG*, 2023.

609 Haolin Li, Jinyang Liu, Mario Sznaier, and Octavia Camps. 3D-HGS: 3D Half-Gaussian Splatting.
610 *ArXiv*, 2024.

611

612 Rong Liu, Dylan Sun, Meida Chen, Yue Wang, and Andrew Feng. Deformable Beta Splatting. In
613 *SIGGRAPH*, 2025.

614 Jie Luo, Tianlun Huang, Weijun Wang, and Wei Feng. A Review of Recent Advances in 3D Gaussian
615 Splatting for Optimization and Reconstruction. *Image and Vision Computing*, pp. 105304, 2024.

616

617 Reza Mahjourian, Martin Wicke, and Anelia Angelova. Unsupervised Learning of Depth and
618 Ego-Motion from Monocular Video Using 3D Geometric Constraints. In *CVPR*, 2018.

619 Zhenxing Mi, Chang Di, and Dan Xu. Generalized Binary Search Network for Highly-Efficient
620 Multi-View Stereo. In *CVPR*, pp. 12991–13000, 2022.

621

622 Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and
623 Ren Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, pp.
624 405–421. Springer, 2020.

625 Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant Neural Graphics
626 Primitives with a Multiresolution Hash Encoding. *ACM TOG*, 2022.

627

628 Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor
629 Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An Imperative Style,
630 High-Performance Deep Learning Library. *NeurIPS*, 2019.

631

632 Johannes L. Schönberger and Jan-Michael Frahm. Structure-From-Motion Revisited. In *CVPR*, 2016.

633

634 Zhuowen Shen, Yuan Liu, Zhang Chen, Zhong Li, Jiepeng Wang, Yongqing Liang, Zhengming Yu,
635 Jingdong Zhang, Yi Xu, Scott Schaefer, Xin Li, and Wenping Wang. SolidGS: Consolidating Gaussian
636 Surfel Splatting for Sparse-View Surface Reconstruction. *arXiv preprint arXiv:2412.15400*,
2024.

637

638 Bin Tan, Rui Yu, Yujun Shen, and Nan Xue. PlanarSplatting: Accurate Planar Surface Reconstruction
in 3 Minutes. In *CVPR*, 2025.

639

640 Christoph Vogel, Konrad Schindler, and Stefan Roth. 3D Scene Flow Estimation with a Piecewise
641 Rigid Scene Model. *IJCV*, 2015.

642

643 Nicolas von Lützow and Matthias Nießner. LinPrim: Linear Primitives for Differentiable Volumetric
Rendering. *arXiv preprint arXiv:2501.16312*, 2025.

644

645 Minye Wu, Haizhao Dai, Kaixin Yao, Tinne Tuytelaars, and Jingyi Yu. BG-Triangle: Bézier Gaussian
646 Triangle for 3D Vectorization and Rendering. In *CVPR*, pp. 16197–16207, 2025.

647

648 Tong Wu, Yu-Jie Yuan, Ling-Xiao Zhang, Jie Yang, Yan-Pei Cao, Ling-Qi Yan, and Lin Gao. Recent
649 Advances in 3D Gaussian Splatting. *Computational Visual Media*, 10(4):613–642, 2024.

648 Jiayu Yang, Jose M. Alvarez, and Miaomiao Liu. Non-Parametric Depth Distribution Modelling
649 Based Depth Inference for Multi-View Stereo. In *CVPR*, pp. 8626–8634, 2022.
650

651 Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. MVSNet: Depth Inference for Unstruc-
652 tured Multi-view Stereo. In *ECCV*, 2018.

653 Zehao Yu, Torsten Sattler, and Andreas Geiger. Gaussian Opacity Fields: Efficient Adaptive Surface
654 Reconstruction in Unbounded Scenes. *ACM Transactions on Graphics (TOG)*, 43(6):1–13, 2024.
655

656 Farhad G Zanjani, Hong Cai, Hanno Ackermann, Leila Mirvakhabova, and Fatih Porikli. Planar
657 Gaussian Splatting. In *WACV*, 2025.

658 Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The Unreasonable
659 Effectiveness of Deep Features as a Perceptual Metric. In *CVPR*, 2018.
660

661 Ziyu Zhang, Binbin Huang, Hanqing Jiang, Liyang Zhou, Xiaojun Xiang, and Shunhan Shen.
662 Quadratic Gaussian Splatting for Efficient and Detailed Surface Reconstruction. In *ICCV*, 2025.
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701

702 SUPPLEMENTAL MATERIAL 703

704 Here we include additional material on background geometry used for computing barycentric coordinates,
705 further implementation details, and additional results.

707 S.1 BACKGROUND 708

709 In this section, we introduce the necessary background geometry used to parameterize a ray-triangle
710 intersection in barycentric coordinates. The *barycentric coordinates* of a point on a triangle $\lambda =$
711 $[\lambda^0, \lambda^1, \lambda^2]^T$, s.t. $\sum_{j=1}^3 \lambda^j = 1$, provide a means for expressing the coordinates of the point as a
712 linear combination of the coordinates of the three vertices, $[V^0, V^1, V^2]$. This parameterization is
713 important for graphics applications to be able to efficiently rasterize triangles onto screen space and
714 interpolate the color of a pixel from each vertex.

715 In rendering, we compute the ray-triangle intersections using,

$$716 P = C + (\hat{r}d) \quad (S.1)$$

717 where $C \in \mathbb{R}^3$ is the camera center, $\hat{r} \in \mathbb{R}^3$ is the unit vector for the ray through pixel i , and $d \in \mathbb{R}$
718 is the depth along the ray from the camera center to the intersection point, computed as follows,

$$719 d = \frac{\hat{n} \cdot \overrightarrow{CB}}{\hat{n} \cdot \hat{r}} \quad (S.2)$$

720 where $B \in \mathbb{R}^3$ is the barycenter of the triangle and $\hat{n} \in \mathbb{R}^3$ is its normal.

721 We compute the barycentric coordinates for point P using the triple products between the normal, a
722 triangle edge, and the vector from each vertex to the point,

$$723 \lambda = \frac{1}{\hat{n} \cdot (\overrightarrow{V_0V_1} \times \overrightarrow{V_0V_2})} \begin{bmatrix} \hat{n} \cdot (\overrightarrow{V_1V_2} \times \overrightarrow{V_1P}) \\ \hat{n} \cdot (\overrightarrow{V_2V_0} \times \overrightarrow{V_2P}) \\ \hat{n} \cdot (\overrightarrow{V_0V_1} \times \overrightarrow{V_0P}) \end{bmatrix} \quad (S.3)$$

724 Intuitively, the contribution of each vertex is proportional to the area of the sub-triangle formed by
725 the intersection point P and the other two vertices of the triangle. This weight becomes larger as P
726 approaches the vertex.

737 S.2 IMPLEMENTATION DETAILS 738

739 In this section, we describe our experimental setting and optimization parameters in detail. To
740 begin optimization, similar to previous work (Huang et al., 2024; Chen et al., 2024a), all geometric
741 supervision is disabled, with optimization only being guided initially by the SSIM loss \mathcal{L}_{ssim} . We
742 enable the normal consistency loss \mathcal{L}_{norm} at iteration 7,000 and enable both the smoothness loss
743 \mathcal{L}_{smooth} and connectivity loss \mathcal{L}_{conn} at iteration 10,000, both chosen empirically. All triangles start
744 with an initial opacity (α) value set to 0.1, with opacity for all primitives being reset every 3,000
745 iterations. We run optimization on the scenes from the DTU and Mip-NeRF 360 datasets for 25,000
746 and 30,000 iterations, respectively. Densification and pruning is run every 250 iterations starting
747 after iteration 2,000. The maximum incenter gradient threshold for densification is set to $7.5e^{-5}$ in
748 all experiments.

749 The learning rates for each respective parameter are set as follows:

- 750 • Spherical Harmonics (Δ): $2.5e^{-3}$
- 751 • Opacity (α): $5e^{-2}$
- 752 • Incenter (μ): $[1.5e^{-4}, 2e^{-6}]$
- 753 • Rotation (R): $1e^{-3}$
- 754 • Scale (s): $4e^{-3}$
- 755 • Diffuse Scalar (σ): $1e^{-3}$

756
757 Table S.1: **Run-time comparison between baseline works and RTS on the DTU dataset (Aanæs
758 et al., 2016).**

759

Method	FPS	Run-Times
2DGS (Huang et al., 2024)	1400	0.41 hr
PGSR (Chen et al., 2024a)	1200	1.05 hr
RTS	400	1.50 hr

760
761
762
763
764

765 where the incenter learning rate follows an exponential decay scheduler starting with $1.5e^{-4}$ and
766 ending with $2e^{-6}$ conditioned on the total number of iterations. We also tested a linear decay
767 scheduler leading to similar results.

768 **The initial diffuse scalar, σ_0 , is inversely proportional to the mean distance, d , between each
769 primitive and the three nearest neighboring primitives,**

770

$$\sigma_0 = \frac{\log\left(\left(\frac{\alpha_0}{\gamma}\right) - 1\right)}{d} \quad (S.4)$$

771
772
773

774 **where α_0 is the initial opacity value and γ corresponds to the minimum αw_σ value necessary
775 for a triangle to be rasterized in the forward pass of the network.**

776
777 **Edge Connectivity Overhead** To compute neighboring edge connections, we construct a single
778 KD-Tree containing the midpoints of all triangle edges which is then queried for each edge once. The
779 connected edge indices are the only structure that is stored during optimization. Since the parameters
780 of the primitives are modified during optimization, we need to recompute the KD-Tree and connected
781 edge indices every 250 iterations (aligned with Adaptive Density Control), which ultimately leads to
782 a minimal overhead for moderate size scenes. To be concrete, building the KD-Tree and computing
783 the neighboring indices takes on average 4 seconds for around 300,000 primitives, which is more
784 than the average number of primitives needed to reconstruct the scenes for the DTU dataset. Since
785 this operation only happens every 250 iterations starting after iteration 10,000 (when the connectivity
786 loss is activated), the overhead of this operation only adds roughly 5 minutes to the optimization
787 process, which is about a 6% increase in runtime. Reconstructing larger scenes, or more precisely,
788 scenes that require more primitives, will naturally demand a larger overhead.

789
790 **Run-Time** The run-times for our approach are roughly 1.5 hours on DTU and 4.5 hours on Mip-
791 NeRF 360, with scenes from DTU and scenes from Mip-NeRF 360 having on average 249,941 and
792 1,275,985 primitives, respectively, using 0.5 resolution for the DTU dataset and 0.25 resolution for
793 the Mip-NeRF 360 dataset. **Please see Table S.1 for a comparison with baselines.**

794
795 **Related Work Extension** In their concurrent work, TriangleSplatting, Held et al. (2025a) also
796 propose a shift towards using triangle primitives as the explicit representation for inference-
797 time optimization of a scene. In contrast to our work, each triangle in their approach is
798 parameterized by a set of 3D points, similar to their previous work in 3D Convex Splatting
799 (3DCS) (Vogel et al., 2015). The primitives in our work are parameterized by an incenter,
800 scalar offsets to each vertex, and a rotation quaternion. Parameterizing the primitives in these
801 separate terms allows for a more selective propagation of gradients with different learning rates,
802 which in our experimentation, showed more favorable convergence toward a higher quality
803 reconstruction. The activation function used in TriangleSplatting also differs from ours, as they
804 use the normalized ReLU of the Signed Distance Function (SDF) for each triangle, whereas we
805 apply a Sigmoid function to the SDF for each triangle. This is similar to the formulation used in
806 3DCS, however, we compute the exact SDF, compared to the approximation used in 3DCS. The
807 primitives in our approach are also more expressive than that of TriangleSplatting. Each vertex
808 encodes a separate color instead of using a single color for an entire triangle. As mentioned
809 in the main paper, a key differentiator between the two works is the optimization criteria, in
which our approach provides an avenue for primitive-to-primitive interactions through our soft
connectivity forces. As shown in Table 1, RTS shows a major increase in geometric accuracy
compared to TriangleSplatting. Both works have similar motivations, in which using triangles

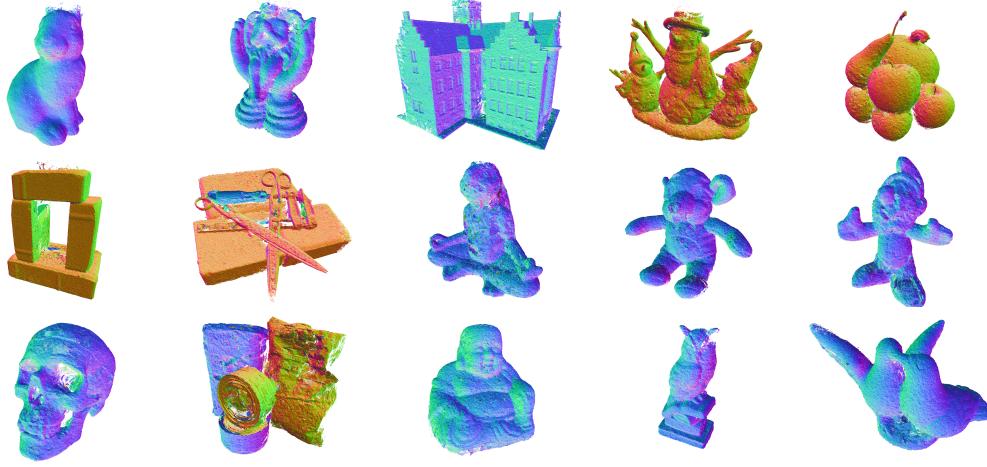


Figure S.1: **Output point cloud visualization of RTS for all scenes from the evaluation set from the DTU dataset(Aanæs et al., 2016).** The points are colored according to their normals.

Table S.2: Novel View Synthesis per-scene on all scenes from the mip-NeRF 360 dataset (Barron et al., 2022).

Scene	SSIM	PSNR	LPIPS
room	0.919	31.25	0.156
counter	0.900	28.27	0.139
kitchen	0.936	30.69	0.083
bonsai	0.930	30.93	0.142
bicycle	0.637	20.31	0.365
flowers	0.617	19.99	0.351
garden	0.742	24.89	0.250
stump	0.658	23.25	0.410
treechill	0.631	18.64	0.371

as the primitive in an alpha-blending optimization framework provides a direct route to the estimation of a mesh with high-quality appearance that can directly be rendered in novel views.

S.3 LOSS FUNCTIONS

For completeness, we define the loss terms used in this paper that were introduced in previous work (Kerbl et al., 2023; Huang et al., 2024). The SSIM loss is computed as follows:

$$\mathcal{L}_{ssim} = \frac{1}{N} \sum_{i,j} (1 - \gamma) L_1 + \gamma L_{D-SSIM} \quad (S.5)$$

The normal consistency loss is computed as follows:

$$\mathcal{L}_{norm} = \frac{1}{N} \sum_{i,j} (1 - \hat{n}^T \hat{n}_d) \quad (S.6)$$

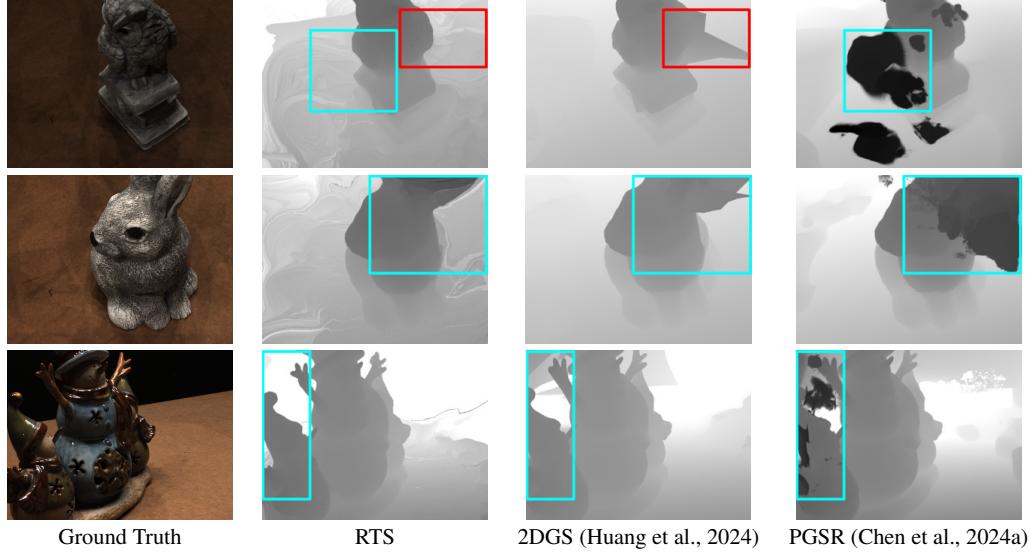
This encourages alignment between the rendered normal \hat{n} with the normal computed from the rendered depth map \hat{n}_d ,

S.4 ADDITIONAL EVALUATIONS

In Table S.2, we show the per-scene novel view synthesis evaluation for RTS on the mip-NeRF 360 dataset (Barron et al., 2022) using the standard metrics of PSNR, SSIM, and LPIPS. As von Lützow & Nießner (2025) note, using primitives with explicit boundaries can begin to introduce hard edges

864
865 Table S.3: Chamfer distance evaluation on scenes from DTU (Aanæs et al., 2016). Following previous
866 literature, we average the accuracy and completeness Chamfer distances on the widely used evaluation
867 set. Chamfer distances are measured in millimeters. **Top:** Chamfer distance evaluation using the GT
868 mask. **Bottom:** Chamfer distance evaluation *without* using the GT mask.

Method	DTU														Mean (mm)↓	
	24	37	40	55	63	65	69	83	97	105	106	110	114	118	122	
2DGGS (Huang et al., 2024)	0.48	0.91	0.39	0.39	1.01	0.83	0.81	1.36	1.27	0.76	0.70	1.40	0.40	0.76	0.52	0.80
PGS (Chen et al., 2024a)	0.36	0.57	0.38	0.33	0.78	0.58	0.50	1.08	0.63	0.59	0.46	0.54	0.30	0.38	0.34	0.52
RTS	0.42	0.61	0.74	0.39	0.53	0.86	0.70	0.84	0.72	0.37	0.68	0.87	0.34	0.58	0.44	0.61
2DGGS (Huang et al., 2024) [no mask]	1.22	1.69	0.88	0.43	0.96	0.77	0.85	1.23	1.87	1.69	0.91	2.01	0.82	0.82	1.07	1.15
PGS (Chen et al., 2024a) [no mask]	1.10	1.59	0.97	0.45	1.98	0.64	0.59	1.70	1.69	1.57	0.74	0.63	0.44	0.71	0.74	1.04
RTS [no mask]	0.52	0.92	0.94	0.55	0.91	1.01	0.74	1.09	1.05	0.67	0.88	1.07	0.42	0.79	0.80	0.82



892 Figure S.2: **Qualitative results between RTS and two baselines on the DTU dataset (Aanæs et al.,**
893 **2016). RTS is much more effective at removing floaters. This is especially helpful in extreme**
894 **viewpoints with low camera overlap.**

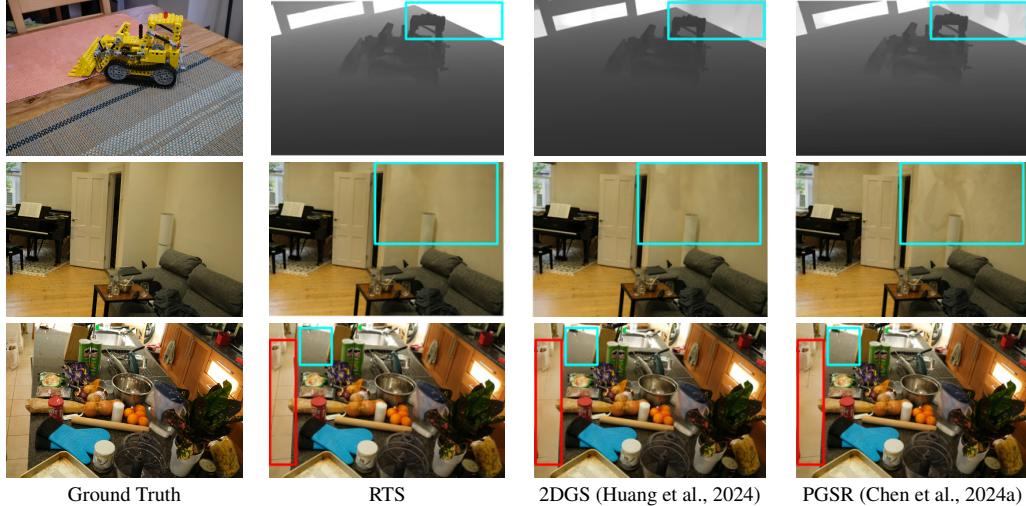
895
896 in regions with poor visibility while smoother primitives degrade more gracefully. While the diffuse
897 boundaries of our triangles help prevent much of this behavior, its effects are noticeable in some of
898 the reconstructed images, especially in outdoor scenes with distant background foliage. See Fig. S.5

900 As mentioned in Section 5.3, we compare the Chamfer distances of RTS and two competitive GS
901 algorithms in Table S.3 with and without the use of the GT segmentation masks. While the distances
902 increase for all methods, this experiment demonstrates how RTS is more effective at floater removal
903 and background modeling.

904 **Qualitative comparisons to competitive baselines on the DTU dataset (Aanæs et al., 2016) are**
905 **shown in Fig. S.2, and on the mip-NeRF 360 dataset (Barron et al., 2022) are shown in Fig. S.3.**
906 Additionally, we provide renderings of novel views on challenging indoor and outdoor scenes with
907 fine details and non-Lambertian surfaces. Output renderings and 3D models for all scenes will be
908 made publicly available to the research community if the paper is accepted. We show qualitative
909 geometric reconstruction results on all scenes from the DTU dataset in Fig. S.1 and novel view
910 synthesis results on all indoor and outdoor scenes from the Mip-NeRF 360 datasets in Fig. S.4 and
911 Fig. S.5, respectively.

912
913
914
915
916
917

918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935



936 **Figure S.3: Qualitative results between RTS and two baselines on the mip-NeRF 360 dataset (Barron et al., 2022). RTS can better represent texture-less areas and recover sharp details from**
937 **surfaces seen in few views.**

938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969



970 **Figure S.4: Visualizations of all indoor scenes, (top-down) [bonsai, counter, kitchen, room], from the**
971 **mip-NeRF 360 dataset (Barron et al., 2022).**

972

973

974

975

976

977

978

979

980

981

982

983

984

985

986

987

988

989

990

991

992

993

994

995

996

997

998

999

1000

1001

1002

1003

1004

1005

1006

1007

1008

1009

1010

1011

1012

1013

1014

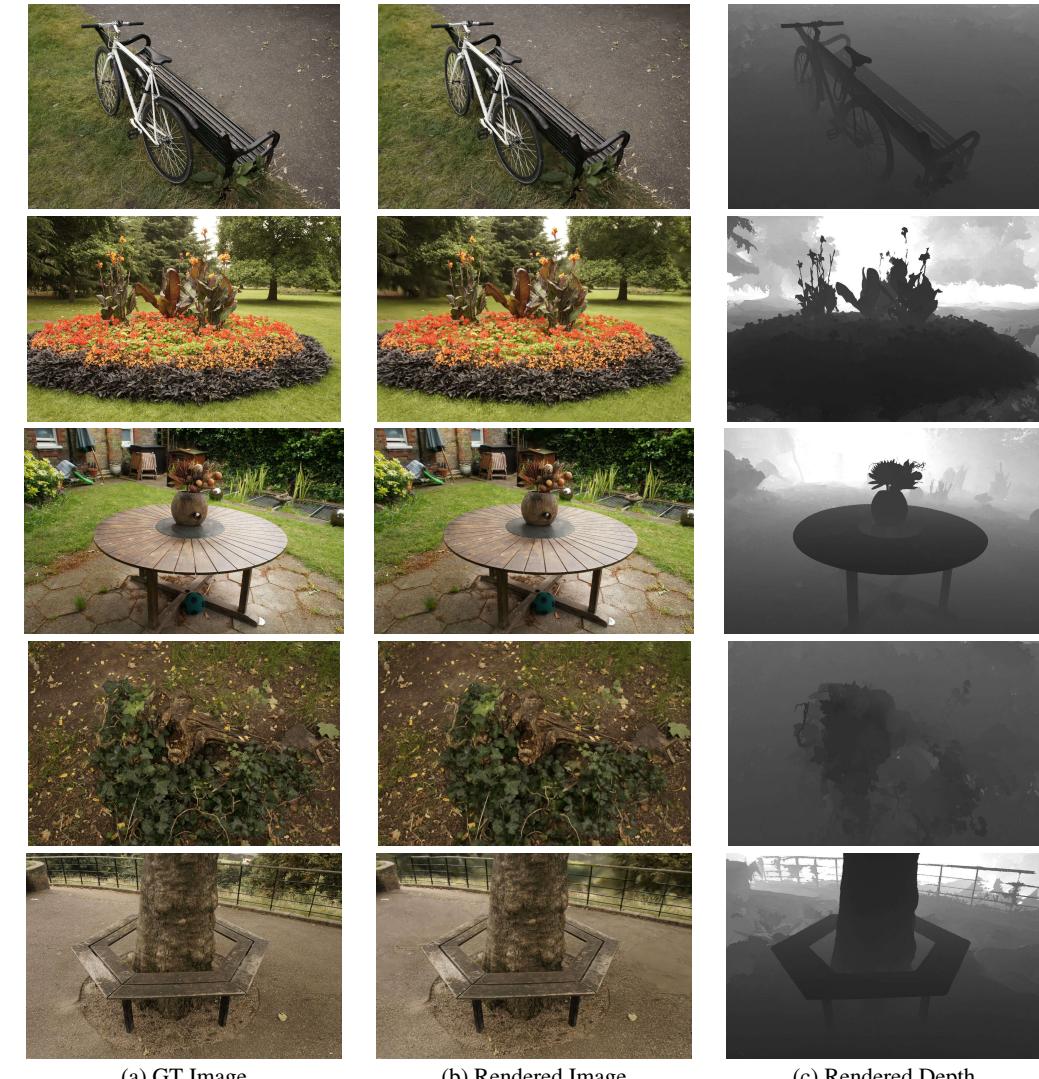


Figure S.5: Visualizations of all outdoor scenes, (top-down) [bicycle, flowers, garden, stump, treehill], from the Mip-NeRF 360 dataset (Barron et al., 2022). In the top-right section of the rendered image of treehill, RTS imprecisely approximates the appearance of the background foliage, while GS-based algorithms typically blur this region.

1019

1020

1021

1022

1023

1024

1025