SDFoam: Signed-Distance Foam for explicit surface reconstruction

Supplementary Material

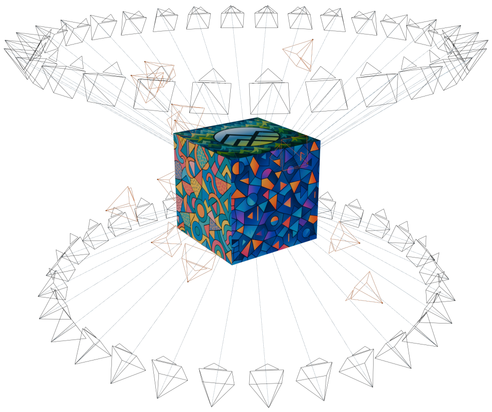

Figure 7. Synthetic scene made with Blender. The cameras placed on the two parallel circumferences (in black) are used as the training set, while the others (in orange) form the test set.

6. Synthetic 3D reconstruction benchmark

To evaluate the ability of our method to recover clean and watertight surfaces, we construct a controlled synthetic scene in Blender consisting of a textured cube observed by 72 calibrated cameras (60 used for training and 12 for testing), as shown in Fig. 7. We train both RadiantFoam (RF) and our SDFoam model using identical hyperparameters and camera configurations to ensure a fair, one-to-one comparison.

This setup exposes a characteristic failure scenario of RF: despite the simplicity of the underlying geometry, the reconstructed density field often develops discontinuities and *holes* even on perfectly planar surfaces (Fig. 12). This issue arises because RF represents density as an independent trainable parameter per Voronoi cell, with no enforced consistency across cell boundaries.

In contrast, SDFoam leverages a continuous signed distance field to impose geometric coherence across the entire scene. Rather than learning density directly, we compute it as the derivative of a sigmoid applied to the SDF, modulated by a learnable sharpness parameter, as defined in Eqs. 5 and 6. This formulation produces a smooth, structurally consistent density field and prevents the surface fragmentation observed in RF.

7. Real-world Scenario

We use DTU as our main benchmark because its accurate ground-truth geometry makes it especially suitable for evaluating mesh quality and surface reconstruction. However, to show that our method is not limited to this setting, we also evaluate it on a different type of dataset with large-scale real-world scenes. In this setting, our performance remains comparable to RadiantFoam, with only a 2-point PSNR gap on the Bonsai scene (30.80 vs. 32.15; Fig. 9).

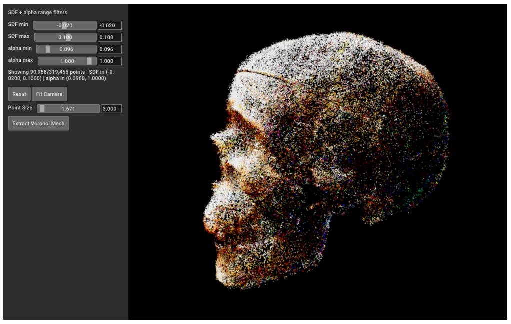

SDF + alpha range filters

SDF min -0.020

SDF max 0.100

alpha min 0.096

alpha max 1.000

Showing 90,958/319,456 points | SDF in (-0.0200, 0.1000) | alpha in (0.0960, 1.0000)

Point Size 1.671 3.000

Figure 8. SDFoam GUI. The voronoi seeds can be filtered using a combination of a SDF threshold and an alpha threshold, and the geometry can be computed by obtaining the corresponding Voronoi vertices and faces from the Voronoi diagram.

Figure 9. Mip-NeRF 360. Left: **SDFoam**, Right: **Ground truth**.

8. Ablation Studies

We perform an ablation study on the DTU dataset to assess the contribution of the different loss components. In particular, we analyze the effect of the point cloud initialization and the Eikonal regularization on both rendering quality and geometric accuracy across three representative scans.

SfM	EIK	Scan 24		Scan 37		Scan 65	
		PSNR ↑	CD ↓	PSNR ↑	CD ↓	PSNR ↑	CD ↓
×	✓	20.82	6.37	22.87	8.42	25.76	7.25
✓	×	22.17	3.30	23.54	3.72	25.78	1.78
✓	✓	**29.80**	**1.86**	**30.42**	**2.87**	**32.43**	**1.52**

Table 4. Ablation study on some DTU scans evaluating the effect of point cloud initialization (SfM) and Eikonal regularization (EIK) on rendering quality (PSNR ↑) and geometric accuracy (Chamfer Distance ↓).

9. Surface flatness results

Figure 12 reports additional qualitative results that highlight the behavior of the two implicit surface representations. We visualize the depth fields produced by RF (RadiantFoam) and by our SDFoam model. As shown, SDFoam is able to recover a consistent and complete geometry even in the presence of complex textures, whereas RF struggles to maintain geometric coherence and often produces holes or incomplete surfaces. This issue is inherited from the original NeRF formulation, where geometry is only indirectly induced and becomes unstable in texture-rich regions. In contrast, the explicit SDF formulation in SDFoam produces more robust and stable reconstructions while preserving a PSNR that remains quantitatively and qualitatively comparable to RF.

10. SDFoam GUI

We developed an interactive GUI (Figure 8) that enables loading a trained scene and visualizing all the Voronoi sites together with their associated colors. The GUI provides fine-grained control over the selection of the sites that contribute to the final geometry. In particular, each site is associated with an SDF value, which we obtain by querying the trained MLP at the site position, and with an alpha value representing its opacity. The interface therefore exposes two independent filtering modules: one operating on the SDF range and one operating on the alpha range.

The user can select a minimum and maximum threshold for both quantities, and a Voronoi site is retained only if it satisfies both conditions simultaneously. This dual filtering mechanism is crucial, since SDF alone is often insufficient to isolate clean geometric structures. By fine-tuning the two thresholds jointly, the user can interactively refine the subset of Voronoi cells that correspond to the actual surface of the reconstructed object. Fig 14 shows an example.

Once a satisfactory subset of sites has been selected, the GUI provides a one-click tool to explicitly compute the full Voronoi diagram restricted to the retained cells. For each site, we compute the corresponding Voronoi region, extract its polygonal faces, and identify all its vertices, producing an explicit polygonal description of the diagram. The resulting mesh is immediately displayed in the GUI's 3D viewport, allowing the user to visually inspect the reconstructed geometry. Additionally, the mesh can be exported as a standard mesh file, as seen in Fig. 13.

Figure 10. Left: RF, Right: SDFoam. Our method successfully gets rid of floaters by leveraging the per-cell SDF values and converting them to density.

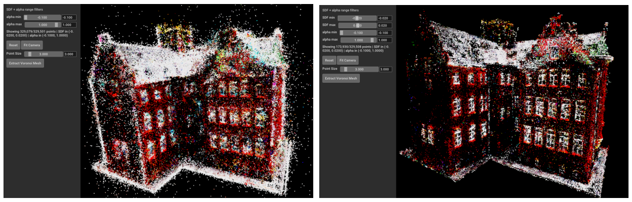

SDF + alpha range filters
alpha min -0.100 -0.100
alpha max 1.000 1.000
Showing 329,079/329,501 points | SDF in (-0.0200, 0.0200) | alpha in (-0.1000, 1.0000)
Point Size 3.000 3.000

SDF + alpha range filters
SDF min -0.020 -0.020
SDF max 0.020 0.020
alpha min -0.100 -0.100
alpha max 1.000 1.000
Showing 173,930/329,508 points | SDF in (-0.0200, 0.0200) | alpha in (-0.1000, 1.0000)
Point Size 3.000 3.000

Figure 11. Left: RF, Right: SDFoam. On RF we can filter out cells based on their alpha value. However, sometimes thresholding is not enough and floating or unwanted sites remain. On SDFoam, we can filter sites both by their alpha value and SDF value, precisely removing any unwanted site. The processed SDFoam scene can be left as is, or converted into a colored point cloud or mesh.

11. Mesh Extraction

We also report qualitative results of the meshes extracted from the SDF field shown in Fig. 12. Figure 11 compares the output of our method with the one extracted from RF. In the case of RF, the mesh is obtained through a naïve procedure. First, we assign a label to each site based on its density value, which allows us to identify the cells belonging to the object. Then, for every cell, we select the faces that are shared with adjacent cells whose density differs significantly. These face discontinuities are then used to extract the final mesh. In other words, we retain only the faces where the density contrast between neighboring cells is high (where the labels differ) and treat these as the polygonal faces of the resulting mesh. In contrast, our method extracts the surface directly from the SDF field by applying a threshold that can be tuned depending on the scene. As shown on the left of Fig. 14, this threshold can be adjusted interactively through the control panel. The extracted geometry already retains the color information: during extraction, each Voronoi cell is assigned the color of its corresponding site. This naturally produces a texture, since the color information is inherently encoded in the Voronoi representation. An example of the extracted mesh of the cube, with the texture applied by SDFoam GUI, is shown on the right of Fig. 14.

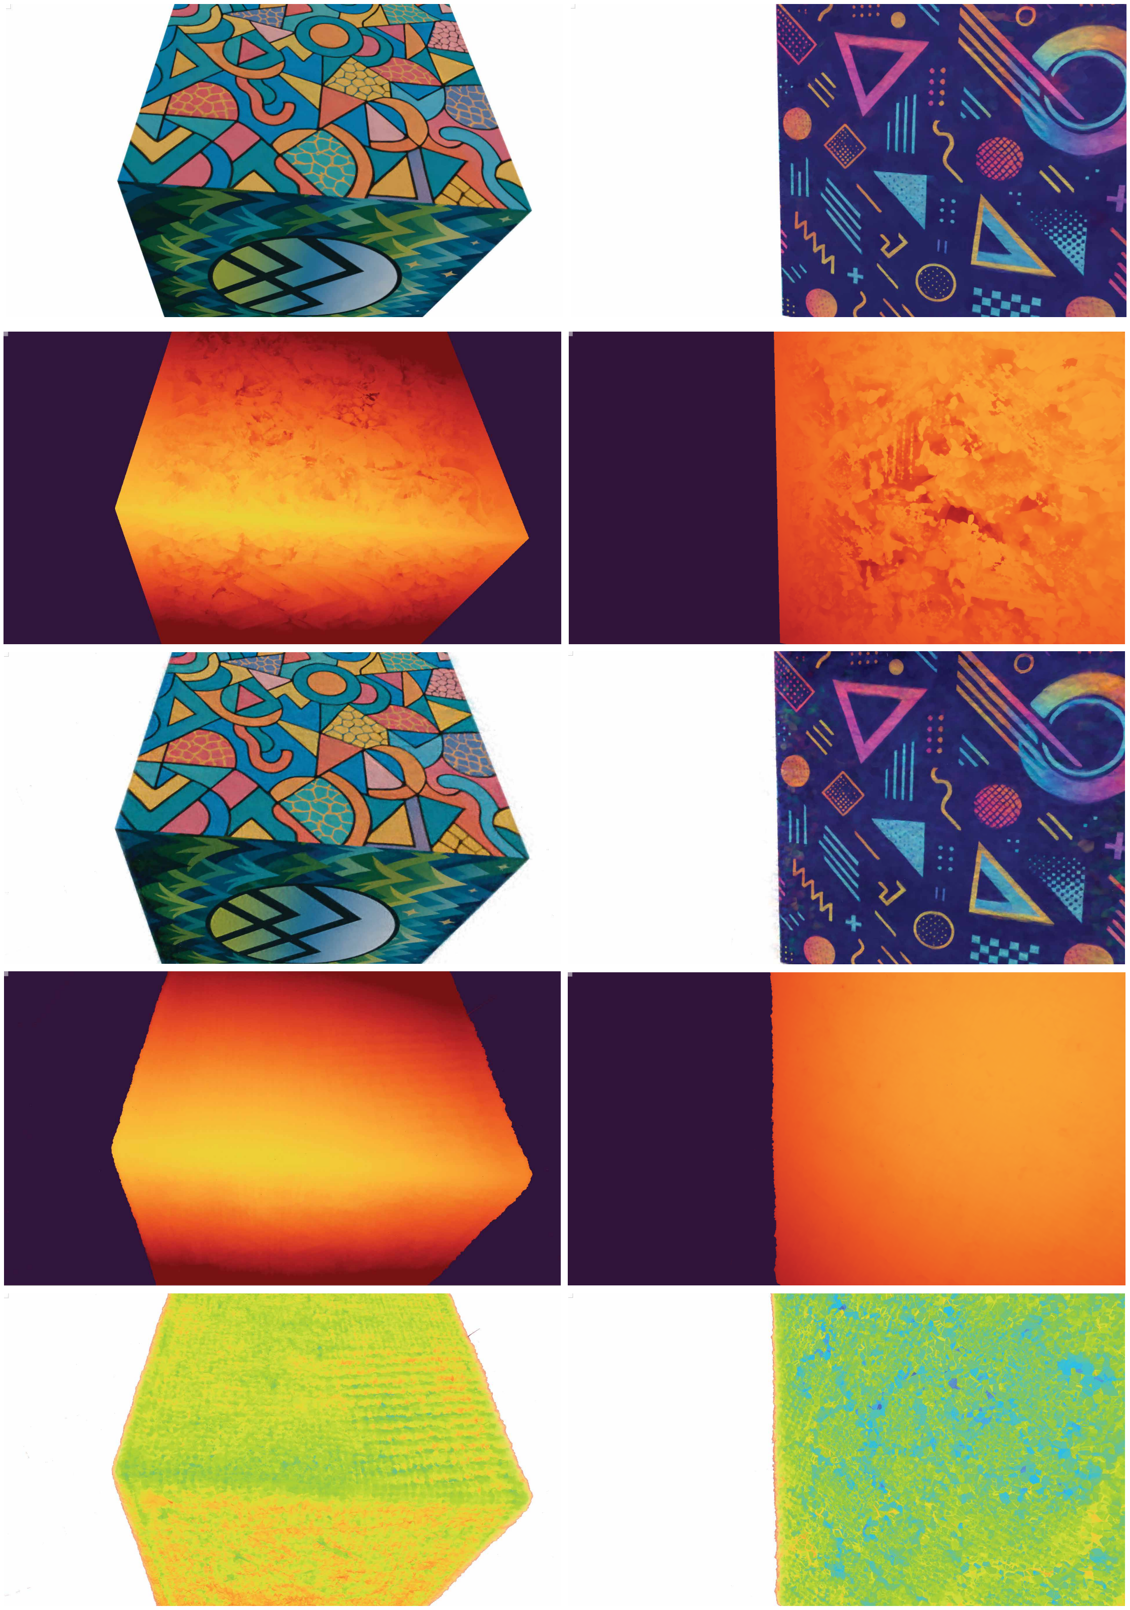

Figure 12. Qualitative comparison of geometry and viewpoint rendering. From top to bottom: RadiantFoam RGB rendering, RadiantFoam depth, SDFoam RGB rendering, SDFoam depth, and (last row) SDFoam per-cell SDF.

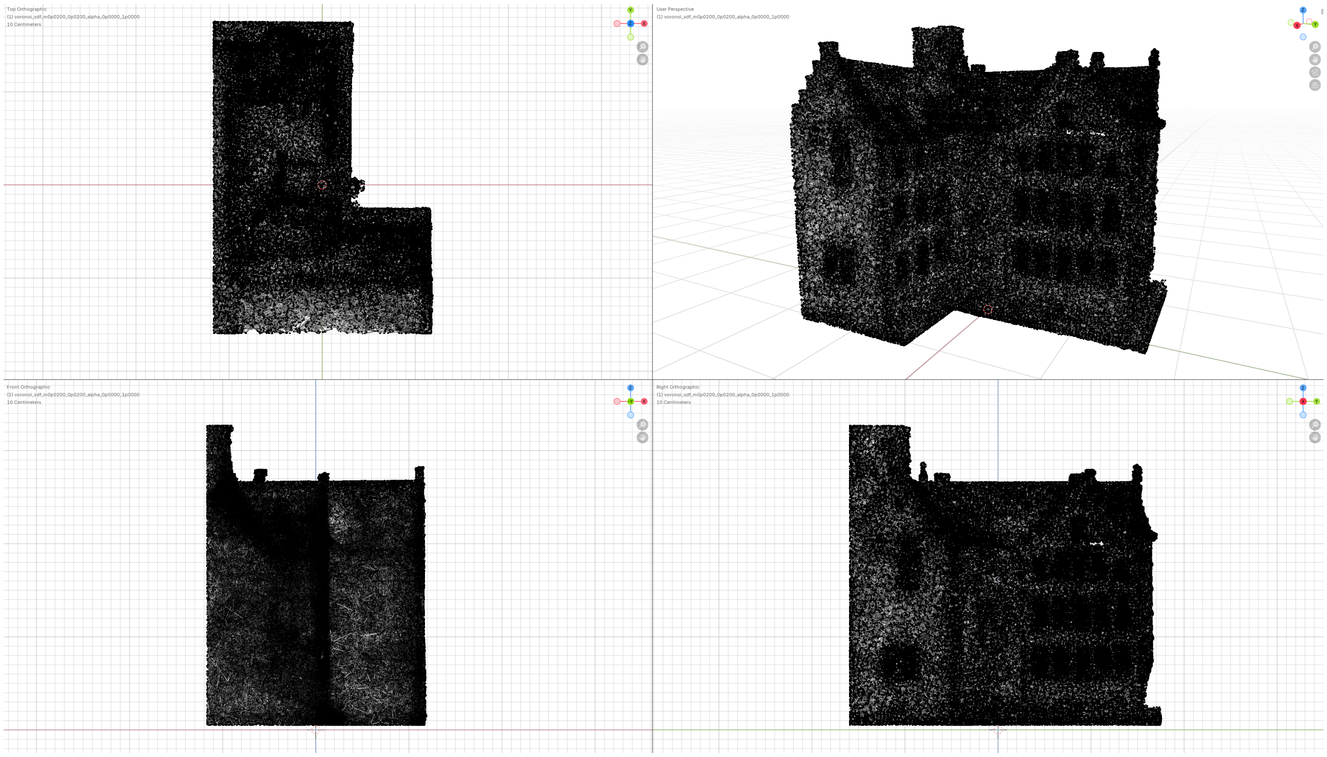

Figure 13. Loading the SDFoam extracted mesh into Blender or any other software for further processing. As seen from the wireframe orthographic views, the walls of the building remain straight, and the high poly count allows each face to retain a sigle color, which is needed for visual fidelity. A lower polygon count mesh can be obtained at this stage through remeshing and uv remapping.

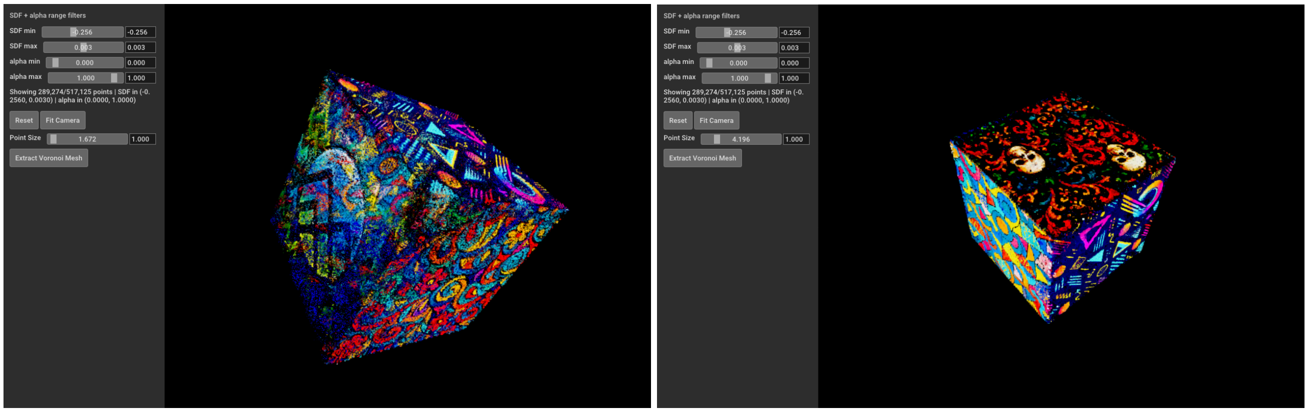

SDF + alpha range filters
SDF min -0.256 -0.256
SDF max 0.003 0.003
alpha min 0.000 0.000
alpha max 1.000 1.000
Showing 289,274/517,125 points | SDF in (-0.2560, 0.0030) | alpha in (0.0000, 1.0000)
Point Size 1.672 1.000

SDF + alpha range filters
SDF min -0.256 -0.256
SDF max 0.003 0.003
alpha min 0.000 0.000
alpha max 1.000 1.000
Showing 289,274/517,125 points | SDF in (-0.2560, 0.0030) | alpha in (0.0000, 1.0000)
Point Size 4.196 1.000

Figure 14. **SDFoam GUI**. Left: filtering the Voronoi sites through alpha + SDF thresholding. Right: the output extracted textured mesh.