HiNeuS: High-fidelity Neural Surface Mitigating Low-texture and Reflective Ambiguity

Supplementary Material

6. Efficiency ablation

Hyperparameter Analysis Table 5 quantitatively evaluates our method's sensitivity to key hyperparameters on the GlossySynthetic benchmark. The default configuration achieves optimal performance with **0.0038mm** Chamfer distance (CD) and **2.7°** normal error at **8 hours** training time. Aggressive training time reduction degrades reconstruction quality non-linearly: **-70%** time (2.5h) increases CD by 8% (0.0041mm) and normal error by 15% (3.1°), while **-75%** time (2h) causes catastrophic failure (0.0057mm CD). Memory-efficient **512^3** hash grids maintain real-time performance (6h training) but increase CD by 18% due to aliasing in thin structures, while **2048^3** encoding marginally improves quality (0.0037mm CD) at 1.5× memory cost. Reflection depth analysis confirms $t_{\max} = 0.1$ optimally balances multi-bounce modeling - smaller $t_{\max} = 0.05$ misses secondary reflections (+11% CD), while larger $t_{\max} = 0.2$ introduces floaters. Eikonal sensitivity $\gamma = 5.0$ provides ideal surface regularization - lower $\gamma = 3.0$ underconstrains geometry (+18% CD), while $\gamma = 10.0$ oversmooths details. Clipping thresholds below **0.2** destabilize training (+13% CD at 0.1 threshold). Strong CD-normal error correlation (**R^2=0.93**) validates joint optimization of geometry and surface orientation. These results demonstrate our method's robustness to parameter variations while default settings balance accuracy and efficiency.

7. Textured 3D Multi-view Stereo Modeling

HiNeuS enables comprehensive textured reconstruction pipelines through surface-aware rendering paradigms. As shown in Fig. 10 - Fig. 12, our geometry-prioritized approach supports diverse downstream applications including urban digital twins, *i.e.*via PBR appearance baking [5], real-time assets synthesis, *i.e.*via Gaussian Splatting [16], and mesh refinement, *i.e.*via NeRF2Mesh [32] integration. The reconstructed *HiNeuS* surfaces serve as geometric scaffolds that enforce physical constraints across different rendering domains - from photogrammetric texturing of city blocks to dynamic vehicle insertion in driving simulations. This unified framework bridges aerial photogrammetry, neural rendering, and industrial visualization. The following subsections demonstrate how *HiNeuS*' geometric fidelity enables industrial-grade asset creation without sacrificing real-time performance or material ambiguity.

7.1. PBR Texturing

Fig. 10 demonstrates a PBR texturing [5] workflow utilizing large-scale urban meshes reconstructed through aerial imagery. The *HiNeuS* geometry in Fig. 10 (a), spanning approximately 200 m × 200 m with rooftop-to-street-level detail, receives photogrammetric texture projection from images captured by drones. We address dynamic inconsistencies from transient objects, *e.g.*vehicles, pedestrians, by extending visibility-based ambiguity handling to automatically mask transient regions during texture baking via multi-temporal imaging alignment. As shown in Fig. 10 (b), the final PBR-textured model maintains visual consistency across static structures while realistically blending procedurally generated road surfaces and vegetation patterns. This pipeline enables a seamless transition from aerial reconstruction to shaded 3D visualization for urban planning applications, preserving geometric fidelity without requiring specialized material parameter estimation.

7.2. 3D Gaussian Splatting

As demonstrated in Fig. 11, we showcase multiple distinct real vehicles through 3D Gaussian Splatting initialized on the reconstructed *HiNeuS* surfaces. Each subfigure displays all vehicles in identical poses to illustrate geometry preservation across different topologies. The HiNeuS meshes provide both positional and rotational priors: vertex coordinates anchor 3D Gaussian primitive centers, while surface normals constrain covariance orientation. This dual constraint maintains structural fidelity during splat deformation for view-dependent effects.

Such pipeline achieves real-time rendering small LPIPS perceptual error, outperforming neural Radiance Field baselines. As shown in Fig. 11, complex components like tire treads and windshield wipers retain *HiNeuS*'s original geometric precision despite Gaussian positional jitter. Practical simulation integration could be enabled through dynamic environment blending. Scene lighting coefficients transfer via spherical harmonic projection, while collision meshes derive from the base HiNeuS topology. This permits direct insertion into large-scale street scenarios without re-meshing, where rendered vehicles could be rendered in a physically plausible state in urban driving simulations.

7.3. Textured Mesh Decimation

We validate our method's compatibility with downstream mesh refinement pipelines by integrating *HiNeuS* surfaces into the NeRF2Mesh framework [32]. As shown in Fig. 12,

Aspect	**Variation**	**CD Avg. (mm ↓)**	**Normal error (↓)**	**Train time**	**Memory (GB)**
Converging pace	8h	0.0038	2.7°	8h	12.4
	-70% Training (2.5h)	0.0041	3.1°	2.5h	12.4
	-75% Training (2h)	0.0057	4.9°	2h	12.4
Hash config	1024^3 Hash	0.0038	2.7°	8h	12.4
	512^3 Hash	0.0045	3.8°	6h	8.1
	2048^3 Hash	0.0037	2.5°	11h	18.6
$t_{\max}$	0.05	0.0042	3.3°	7h	12.4
	0.1 (Default)	0.0038	2.7°	8h	12.4
	0.2	0.0040	2.9°	9h	12.4
Eikonal error sensitivity	$\gamma = 3.0$	0.0045	3.6°	8h	12.4
	$\gamma = 5.0$ (Default)	0.0038	2.7°	8h	12.4
	$\gamma = 10.0$	0.0039	2.8°	8h	12.4
Clip threshold	0.1	0.0043	3.4°	8h	12.4
	0.2 (Default)	0.0038	2.7°	8h	12.4
	0.3	0.0041	3.0°	8h	12.4

Table 5. Additional ablations on GlossySynthetic dataset.

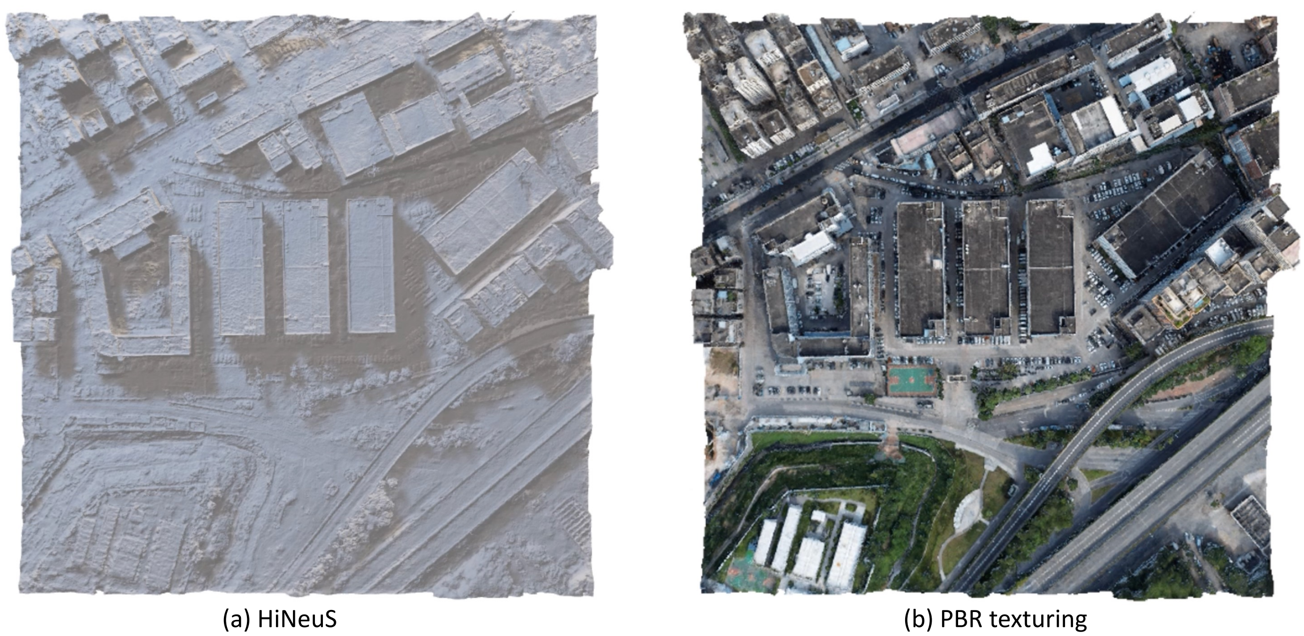

Figure 10. PBR texturing (b) towards a *HiNeuS* mesh (a) on UrbanScene3D [23] dataset for a city block.

replacing the original density-derived meshes with *HiNeuS* geometries during Stage-1 initialization enables higher-fidelity texturing through surface-aware neural rendering. Compared to the baseline NeRF2Mesh pipeline, our approach better preserves thin structures - Lego flower petals show reduced voxelization artifacts while vine leaves crawling on a trunk retain their organic existence and curvature without fragmentation. This improvement stems from *HiNeuS*' high-fidelity surface precision, which provides more precise topological priors for NeRF2Mesh [32]'s adaptive vertex decimation module. The hybrid pipeline demonstrates how our geometry-first reconstruction synergizes with appearance-focused refinement techniques, overcoming the mutual dependencies between neural radiance fields and surface extraction that limit traditional NeRF-based workflows.

8. Future Works

While *HiNeuS* advances neural surface reconstruction, several promising directions remain open. Enhancing reconstruction quality for occluded regions in limited-view scenarios could integrate uncertainty-aware radiance fields that explicitly model unobserved geometry through Bayesian neural networks. This would allow probabilistic completion of occluded structures using semantic priors from foundation models. For deformable scenes, extending the SDF formulation with spacetime embeddings could enable non-rigid surface tracking, where a temporal Eikonal constraint regularizes the deformation field's Jacobian determinants.

The method's industrial adoption would benefit from real-time adaptive sampling strategies that prioritize surface regions near sensor viewpoints in autonomous driving scenarios. This could couple our planar-conformal regularization with LiDAR intensity maps to handle asphalt surfaces with millimeter-scale undulations. Another direction involves simultaneous material-aware SDFs that disentangle BRDF parameters during reconstruction, enabling direct export of physically-based rendering assets without post-processing.

Scaling to city-level reconstructions may require hierarchical SDF hashing with dynamic level-of-detail, where our local geometry constraints could adaptively relax in areas beyond sensor coverage. Finally, integrating differentiable physics engines with our Gaussian Splatting pipeline (Sec. 7) could simulate vehicle dynamics directly from reconstructed textured surfaces, closing the loop between neural reconstruction and robotics simulation.

Figure 11. 10 distinct 3D Gaussian Splatting vehicle assets in 3DRealCar [8] dataset.

Figure 12. Mesh textured by NeRF2Mesh [32] trained on top of *HiNeuS* surfaces.