

Neural Gabor Splatting: Enhanced Gaussian Splatting with Neural Gabor for High-frequency Surface Reconstruction

Supplementary Material

1. Experiment Details

To ensure a fair comparison across different methods, we constrained the number of primitives such that the total data size remained roughly the same. To make the data size similar, we defined the initial number of primitives and the maximum allowable number of them for each dataset and method. Tables 1 and 2 summarize these settings. Across all datasets, the number of primitives after training stays within 100–200 of the designated maximum. This small discrepancy arises from the interplay between densification and pruning.

Densification and Pruning. Densification was performed every 100 iterations. For 3DGS and 2DGS, we adopt the default densification thresholds from the original implementations. Densification continued until 15,000 iterations, after which further accuracy gains saturated and the primitive count had nearly reached the predefined maximum. Once the predefined maximum was reached, new primitives were added only from those having the largest position gradients, ensuring that the primitive number does not exceed the limit. The pruning follows the default strategy.

For 3D Gabor splatting, achieving perfect fairness in data size is difficult because the method does not model view-dependent appearance. Therefore, instead of matching the total data size, we matched the *number of primitives* to our method to ensure a balanced and interpretable comparison. Densification and pruning followed the same scheduling as our method.

Frequency-Aware Densification. For our method, frequency-aware densification was applied for 19,000 out of the 20,000 training iterations. During each densification step, we randomly sampled 20 training views, processed them as a batch, and applied FFT and inverse FFT filtering. The per-pixel frequency-domain error from each view was projected onto the corresponding primitives, and for each primitive, the maximum error over the 20 views was used as the densification score.

Datasets. For the DTU dataset, we used scenes `scan83`, `scan105`, `scan106`, and `scan114`. For Tanks and Temples, we evaluated the `train` and `truck` scenes. Across all datasets, we used a 1/8 split for testing, with the remaining images used for training.

2. Additional Quantitative Results

We further compare our method with Textured Gaussians [1], which achieves high-fidelity reconstruction by assigning an individual texture map to each primitive. Table 3 reports the quantitative results on the High-Frequency dataset [5].

Across all metrics, our method outperforms textured Gaussians under the limited primitive budget. Although textured Gaussians improves appearance quality by training textured primitives after a pre-trained Gaussian splatting stage, this two-stage pipeline implicitly requires a sufficient number of primitives to cover fine-scale structures. When the primitive count is constrained, its pre-training stage tends to overfit high-frequency regions with overspecialized and elongated primitives, reducing the effective representational capacity of the subsequent texture optimization.

In contrast, our method trains the neural Gabor primitives in an end-to-end manner, enabling each primitive to represent diverse color patterns and avoiding degenerate fitting behaviors in high-frequency regions (Fig. 1). As a result, our model preserves complex appearance details even with substantially fewer primitives.

To ensure a fair comparison in terms of data size, we configured Textured Gaussians to use RGBA textures with a resolution of 4×4 , and we set the number of primitives to 10,200. This setting matches the total data footprint of our method while respecting the constraints of the original Textured Gaussians formulation.

The quantitative results for our method in Table 3 match those presented in the main paper.

3. Additional Qualitative Results

We present additional qualitative novel-view synthesis results on the High-Frequency dataset, the Mip-NeRF 360 dataset, and the DTU dataset. We compare our method

Table 1. **Initial and maximum primitive counts for each Mip-NeRF360 scene.** Left value indicates the initial primitive count; right value indicates the maximum count.

Mip-NeRF 360	3DGS	2DGS	Ours	3D Gabor
Bicycle, Garden, Stump	740k-3,700k	740k-3,700k	600k-3,000k	600k-3,000k
Bonsai, Counter, Kitchen, Room	50k-243k	50k-243k	40k-200k	40k-200k

Table 2. **Initial and maximum primitive counts for High-Frequency [5], DTU [2], and Tanks and Temples (tandt) [3].** Left value indicates the initial primitive count; right value indicates the maximum count.

Dataset	3DGS	2DGS	Ours	3D Gabor
High-Frequency	6k-24k	6k-24k	5k-20k	5k-20k
DTU	2.5k-24k	2.5k-24k	2k-20k	2k-20k
train (tandt)	60k-356.3k	60k-356.3k	60k-300k	60k-300k
truck (tandt)	20k-118.8k	20k-118.8k	20k-100k	20k-100k

Table 3. **Additional quantitative comparison between Textured Gaussians [1] and our method on the High-Frequency dataset.** Our method achieves consistently higher performance under the constrained primitive setting.

	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
Textured Gaussians	23.57	0.7912	0.3203
Ours	26.49	0.8808	0.2115

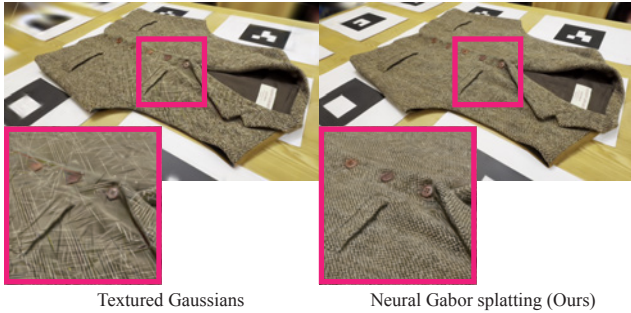


Figure 1. **Visual comparison between Textured Gaussians and our method on High-Frequency dataset.** Our method expresses high-frequency texture successfully, while Textured Gaussians produces elongated monotonous color primitives due to the separation of initial primitive shape training and subsequent texture training.

with 2D Gaussian splatting (2DGS) and 3D Gabor splatting. For fairness, we augment 3D Gabor splatting with our frequency-aware densification strategy. **For all methods, the maximum number of primitives is constrained to match the same overall data size.**

As shown in Fig. 2, our method produces sharper and more faithful reconstructions than both 2DGS and 3D Gabor Splatting on the Mip-NeRF 360 dataset. In particular,

3D Gabor Splatting—while capable of representing high-frequency components—exhibits noticeable color artifacts due to its lack of view-dependent modeling.

Additional comparisons on the High-Frequency and DTU datasets are presented in Fig. 3. Even under strict limitations on the number of primitives, our method successfully preserves fine-grained patterns—such as intricate clothing textures and thin structures—while competing methods suffer from blurring or loss of detail.

Moreover, the supplemental novel-view synthesis videos further demonstrate that our method reconstructs 3D scenes without structural collapse even under view-dependent appearance changes, highlighting the stability of our representation across continuous camera movements.

4. Visual Comparisons with Neural Splatting Methods

We provide additional visual comparisons with NEST [6] and NTS [4] under the same experimental settings as Sec. 5.6. Fig. 4 shows results on the BOOTS scene from the High-Frequency dataset and the BONSAI scene from Mip-NeRF360, each with corresponding zoomed-in views.

We observe that NEST tends to exhibit noticeable visual degradation under tight memory budgets, often leading to structural inconsistencies or artifacts. NTS is affected by suboptimal primitive configurations such as skinny Gaussians, resulting in insufficient fitting that cannot be fully compensated for by its neural texture representation.

In contrast, our method maintains more stable geometry and appearance across both scenes, with fewer visible artifacts under the same budget constraints.




Figure 2. **Additional qualitative comparisons on Mip-NeRF 360 with identical data size.** Our method produces higher visual fidelity compared to 2DGS and 3D Gabor Splatting. 3D Gabor Splatting tends to generate color artifacts (circled) due to the absence of view-dependent modeling.

5. Rendering Time

We report rendering speed (FPS) under the same experimental setup as Sec. 5.8 of main paper (Tab. 4). Due to per-primitive MLP evaluation, our method is slower than 3DGS. However, it achieves significantly higher rendering speed than NEST and NTS, which rely on additional tri-plane or hash-grid computations. This indicates that our method introduces moderate overhead while remaining efficient among neural splatting approaches.

Table 4. Rendering time (FPS) on the BOOTS model under different data budget.

Data Budget	20%	5%	1%
3DGS	242.8	349.7	467.3
NEST	12.25	24.55	43.86
NTS	11.69	15.01	51.89
Ours	85.32	116.8	145.4

References

- [1] Brian Chao, Hung-Yu Tseng, Lorenzo Porzi, Chen Gao, Tuo-tuo Li, Qinbo Li, Ayush Saraf, Jia-Bin Huang, Johannes Kopf, Gordon Wetzstein, et al. Textured gaussians for enhanced 3d scene appearance modeling. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 8964–8974, 2025.
- [2] Rasmus Jensen, Anders Dahl, George Vogiatzis, Engin Tola, and Henrik Aanæs. Large scale multi-view stereopsis evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 406–413, 2014.
- [3] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics (ToG)*, 36(4): 1–13, 2017.
- [4] Yiming Wang, Shaofei Wang, Marko Mihajlovic, and Siyu Tang. Neural texture splatting: Expressive 3d gaussian splatting for view synthesis, geometry, and dynamic reconstruction. In *Proceedings of the SIGGRAPH Asia 2025 Conference Papers*, pages 1–12, 2025.




Figure 3. **Additional qualitative comparisons on the High-Frequency and DTU datasets under identical data size.** Our method better preserves fine-grained patterns—such as intricate clothing textures and thin structures—especially when the primitive count is limited.

- [5] Haato Watanabe, Kenji Tojo, and Nobuyuki Umetani. 3D Gabor Splatting: Reconstruction of High-frequency Surface Texture using Gabor Noise. In *Eurographics 2025 - Short Papers*. The Eurographics Association, 2025.
- [6] Xin Zhang, Anpei Chen, Jincheng Xiong, Pinxuan Dai, Yujun Shen, and Weiwei Xu. Neural shell texture splatting: More details and fewer primitives. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 25229–25238, 2025.




Figure 4. **Visual comparisons with NEST and NTS under matched memory budgets.** Results on the BOOTS scene from the High-Frequency dataset and the BONSAI scene from Mip-NeRF360 are shown. Each group includes the full image and a corresponding zoomed-in image highlighting fine-scale structures.