MAGS: MASK AUGMENTED GAUSSIAN SPLATTING FOR REFLECTIVE SCENES

Anonymous authors
Paper under double-blind review

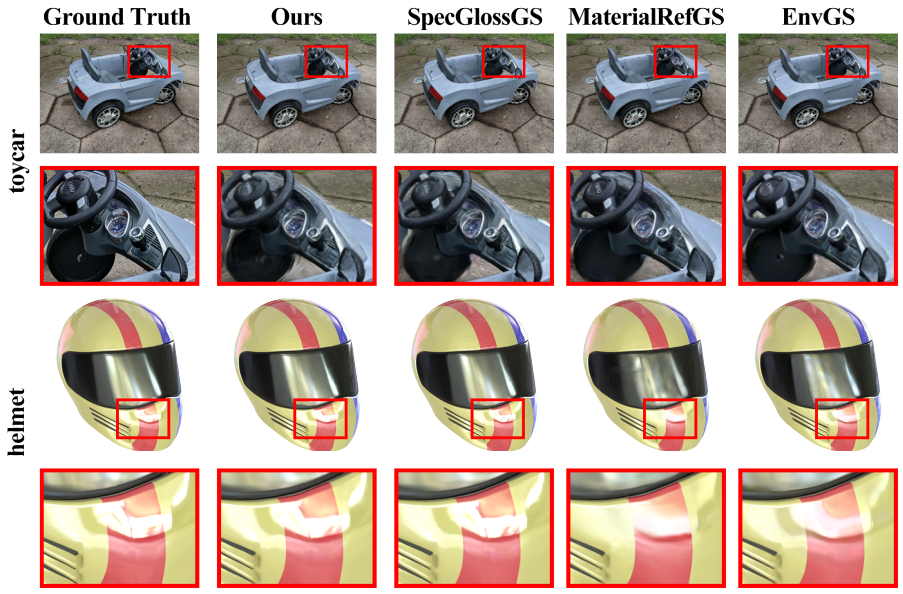


Figure 1: We highlight our novel view synthesis results on scenes with complex reflection effects. MAGS achieves performance competitive with state-of-the-art methods, with best or joint-best quantitative results for most scene–metric pairs on the synthetic ShinyBlender benchmark (Verbin et al., 2024).

ABSTRACT

Novel view synthesis of reflective scenes remains challenging because different combinations of geometry, materials, and illumination can explain similar observed appearances. This ambiguity can lead to inconsistent material estimates within regions that share a common material, potentially degrading reflection rendering at novel viewpoints. We introduce Mask-Augmented Gaussian Splatting (MAGS), a framework that improves reflective novel view synthesis through mask-guided material regularization. Building on a physically based 2D Gaussian Splatting framework, MAGS uses segmentation masks as regional priors and penalizes deviations of rendered material attributes from their mask-level means. The resulting gradients propagate through differentiable rasterization to the underlying Gaussian material parameters, encouraging regional material consistency while allowing rendered appearance to vary across viewpoints. Experiments demonstrate competitive performance on the real-world Ref-Real benchmark and the best average PSNR, SSIM, and LPIPS among the compared methods on the synthetic ShinyBlender benchmark.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

1 INTRODUCTION

Novel view synthesis (NVS) aims to generate photorealistic images from previously unseen viewpoints given a set of images captured from multiple views of a scene (Mildenhall et al., 2019; Barron et al., 2022). Neural Radiance Fields (NeRF) (Mildenhall et al., 2021) established a major paradigm for this task by representing a scene as a continuous volumetric radiance field parameterized by a multilayer perceptron (MLP). NeRF predicts volume density from spatial position and radiance from both position and viewing direction, synthesizing images through differentiable volume rendering. NeRF and its extensions have demonstrated high-quality novel view synthesis of scenes with complex geometry and view-dependent appearance (Mildenhall et al., 2021; Barron et al., 2022). However, the original MLP-based formulation requires repeated network evaluations at sampled points along camera rays, resulting in substantial training and rendering costs (Yu et al., 2021; Fridovich-Keil et al., 2022).

3D Gaussian Splatting (3DGS) (Kerbl et al., 2023) advanced efficient NVS by combining an explicit scene representation with real-time rendering. Unlike the implicit MLP-based representation of the original NeRF, 3DGS represents a scene as a set of optimizable anisotropic 3D Gaussians and renders them through visibility-aware splatting. Direct optimization of these primitives, together with a differentiable rasterizer that avoids repeated neural network evaluations along camera rays, enables high-quality novel view synthesis with efficient training and real-time rendering. This combination of visual quality and computational efficiency has provided a foundation for subsequent advances in Gaussian-based NVS (Yu et al., 2024; Huang et al., 2024; Xie et al., 2025).

Despite these advances, the low-order spherical harmonics (SH) used in conventional Gaussian Splatting have limited capacity to capture sharp specular highlights and high-frequency reflection details (Jiang et al., 2024; Xie et al., 2025). To improve novel view synthesis of reflective scenes, recent methods augment Gaussian primitives with material attributes and incorporate physically based rendering (PBR) to model view-dependent reflections through interactions between surface materials and illumination (Yao et al., 2025; Zhang et al., 2026). SpecGlossGS (Kouros et al., 2026) combines 2D Gaussian Splatting (2DGS) with a specular-glossiness material parameterization and deferred shading. It further uses diffusion-based priors for surface normals and diffuse color to mitigate ambiguity among geometry, materials, and illumination.

Nevertheless, accurately reproducing the observed images does not uniquely determine material properties, as different combinations of materials and illumination can yield similar appearances (Kouros et al., 2024; 2026). Such ambiguity can result in spatially inconsistent material estimates even within surface regions that share a common material (Kouros et al., 2024). In Gaussian-based deferred shading, these inconsistencies can affect the rendered material maps as the primitives' blending weights vary across views, potentially degrading reflection quality at novel viewpoints (Zhang et al., 2026). This motivates encouraging material consistency among Gaussians within regions expected to share a common material as an additional constraint for improving novel view synthesis of reflective scenes.

We introduce Mask-Augmented Gaussian Splatting (MAGS), a framework that improves novel view synthesis of reflective scenes through mask-guided material regularization. Building on the 2DGS-based rendering framework of Kouros et al. (2026), MAGS uses segmentation masks as regional priors to encourage material consistency among Gaussians within each segmented region. The regularized material properties are used in physically based rendering to model view-dependent reflections.

Our contributions are summarized as follows:

- We propose MAGS, a Gaussian Splatting framework for reflective novel view synthesis that combines physically based rendering with regional material consistency constraints.
- We introduce a mask-guided material regularization strategy that uses segmentation masks as regional priors to penalize inconsistencies in material attributes among Gaussians within each region.
- We demonstrate that MAGS achieves quantitative novel view synthesis performance competitive with state-of-the-art methods and improved visual quality on the evaluated real-world and synthetic benchmarks of reflective scenes.

2 Related Work

3DGS (Kerbl et al., 2023) enables real-time novel view synthesis using anisotropic 3D Gaussians and visibility-aware splatting. However, its radiance-based appearance representation does not explicitly separate material from illumination, and its volumetric primitives do not inherently provide surface-aligned geometry and normals for physically based shading. To address these limitations, recent methods extend Gaussian Splatting to support inverse rendering, enabling the estimation of geometry, materials, and illumination from multi-view images. GS-IR (Liang et al., 2024) tackles this problem under unknown lighting by augmenting Gaussians with normal and material attributes and adopting a differentiable physically based rendering pipeline. It employs depth-derivation-based regularization for normal estimation and baking-based occlusion to model indirect illumination. GlossyGS (Lai et al., 2025) focuses on glossy objects, where inverse rendering ambiguity hinders the recovery of plausible geometry and materials. It combines a hybrid explicit–implicit geometry and material representation with normal map prefiltering and a micro-facet geometry segmentation prior. These components improve the modeling of surface normal distributions and reduce ambiguity between geometry and material properties.

2D Gaussian Splatting (2DGS) (Huang et al., 2024) provides a surface-oriented alternative to volumetric Gaussian representations by modeling scenes with oriented planar Gaussian disks. It employs perspective-correct splatting based on ray–splat intersection, together with depth distortion and normal consistency regularization, to improve geometric consistency and surface reconstruction. Because each disk's orientation explicitly defines its surface normal, we adopt this representation as the geometric foundation for our physically based rendering framework. However, the original 2DGS still models view-dependent appearance using spherical harmonics without explicitly separating material from illumination.

Recent approaches have extended the surface-oriented representation of 2DGS to model complex reflective appearance. EnvGS (Xie et al., 2025) addresses the difficulty of capturing near-field reflections and high-frequency reflection details with conventional environment maps. It introduces environment Gaussians as an explicit 3D representation of the reflected environment, alongside base Gaussians for scene geometry and base appearance. A differentiable ray-tracing renderer evaluates the environment Gaussians along reflection directions, and the resulting colors are blended with the base appearance. This approach removes the distant-lighting assumption of conventional environment maps and enables detailed near-field and high-frequency reflections while maintaining real-time rendering. MaterialRefGS (Zhang et al., 2026) addresses under-constrained material inference in deferred shading-based physically based rendering. Because Gaussian contributions vary across viewpoints, rendered material maps can become inconsistent despite view-independent per-Gaussian attributes. It aligns corresponding surface patches using rendered depth and normals to enforce multi-view consistency on diffuse, roughness, and metallic maps. Photometric variations across views further provide a reflection-strength prior, which guides metallic supervision in highly reflective regions using SAM2 masks and regional mean metallic values. Its environment model combines environment-map queries with differentiable 2DGS ray tracing to account for indirect illumination and inter-object occlusion.

For glossy object reconstruction and relighting, Kouros et al. (2026) integrate a microfacet BRDF with specular-glossiness parameterization into 2DGS using deferred shading. This parameterization represents diffuse albedo and specular reflectance as separate learnable attributes. The framework also incorporates diffusion-based priors for surface normals and diffuse color to reduce ambiguity among geometry, materials, and illumination. These studies have advanced Gaussian-based reflection modeling through environment modeling, material estimation, multi-view consistency constraints, and learned priors. MAGS builds on the rendering framework of Kouros et al. (2026) and introduces mask-guided regularization to encourage material consistency among Gaussians within segmented regions, improving novel view synthesis in reflective scenes.

3 Method

MAGS extends the physically based 2DGS rendering framework of SpecGloss-GS (Kouros et al., 2026) with mask-guided material regularization to improve novel view synthesis of reflective scenes.

We first review the underlying Gaussian representations, then describe the rendering framework and the proposed regional material constraints.

3.1 PRELIMINARIES

3D Gaussian Splatting (3DGS) represents a scene as a collection of anisotropic 3D Gaussian primitives. Each primitive is defined as

$$\mathcal{G}_i = \{\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i, \alpha_i, \mathbf{f}_i\}, \tag{1}$$

where $\boldsymbol{\mu}_i \in \mathbb{R}^3$ denotes the Gaussian center, $\boldsymbol{\Sigma}_i \in \mathbb{R}^{3\times3}$ is its covariance matrix, α_i represents opacity, and $\mathbf{f}_i$ denotes the associated appearance feature.

To ensure a valid covariance matrix while enabling efficient optimization, $\boldsymbol{\Sigma}_i$ is parameterized by rotation matrix $\mathbf{R}_i$ and scaling matrix $\mathbf{S}_i$ as

$$\boldsymbol{\Sigma}_i = \mathbf{R}_i \mathbf{S}_i \mathbf{S}_i^\top \mathbf{R}_i^\top. \tag{2}$$

Given a camera view, each 3D Gaussian is projected onto the image plane. The projected Gaussians are then depth-ordered and composited using alpha blending. Specifically, the rendered feature at a pixel $\mathbf{p}$ can be expressed as

$$\mathbf{F}(\mathbf{p}) = \sum_i T_i(\mathbf{p}) \alpha_i(\mathbf{p}) \mathbf{f}_i, \tag{3}$$

where

$$T_i(\mathbf{p}) = \prod_{j<i} (1 - \alpha_j(\mathbf{p})) \tag{4}$$

denotes the accumulated transmittance of the Gaussians preceding the i-th Gaussian along the viewing direction.

2D Gaussian Splatting (2DGS) (Huang et al., 2024) represents a scene as a collection of oriented planar Gaussian disks embedded in 3D space, thereby alleviating the multi-view geometric inconsistency of volumetric 3D Gaussians from 3DGS. Each primitive is defined on a local tangent plane and parameterized as

$$\mathcal{G}_i = \{\boldsymbol{\mu}_i, \mathbf{t}_{i,1}, \mathbf{t}_{i,2}, s_{i,1}, s_{i,2}, \alpha_i, \mathbf{f}_i\}, \tag{5}$$

where $\boldsymbol{\mu}_i \in \mathbb{R}^3$ denotes the center, $\mathbf{t}_{i,1}, \mathbf{t}_{i,2} \in \mathbb{R}^3$ are the two orthonormal unit tangent directions, $s_{i,1}$ and $s_{i,2}$ are the corresponding scale factors, α_i represents opacity, and $\mathbf{f}_i$ denotes the associated appearance feature.

A point on the surface represented by the i-th Gaussian can be parameterized as

$$\mathbf{x}_i(u, v) = \boldsymbol{\mu}_i + s_{i,1} u \mathbf{t}_{i,1} + s_{i,2} v \mathbf{t}_{i,2}, \tag{6}$$

where (u, v) are local coordinates on the tangent plane. The corresponding surface normal is given by

$$\mathbf{n}_i = \mathbf{t}_{i,1} \times \mathbf{t}_{i,2}. \tag{7}$$

Thus, each primitive explicitly defines a surface element together with its local tangent frame and surface normal.

3.2 MASK AUGMENTED GAUSSIAN SPLATTING

To address the material inconsistency problem, we propose MAGS, which encourages material values to remain consistent within each predefined mask region.

Framework Overview. Figure 2 illustrates the overall pipeline of MAGS. We adopt 2DGS as the geometric backbone because its oriented surface primitives provide explicit normals for physically based shading. Building on the rendering framework of SpecGloss-GS (Kouros et al., 2026), we optimize a set of 2D Gaussians with learnable geometric and material attributes from posed input images. These attributes are rasterized into a screen-space G-buffer containing diffuse albedo, roughness, specular reflectance (F_0), indirect illumination, and surface normals. A differentiable deferred renderer combines these quantities with learned environment lighting to compute the diffuse

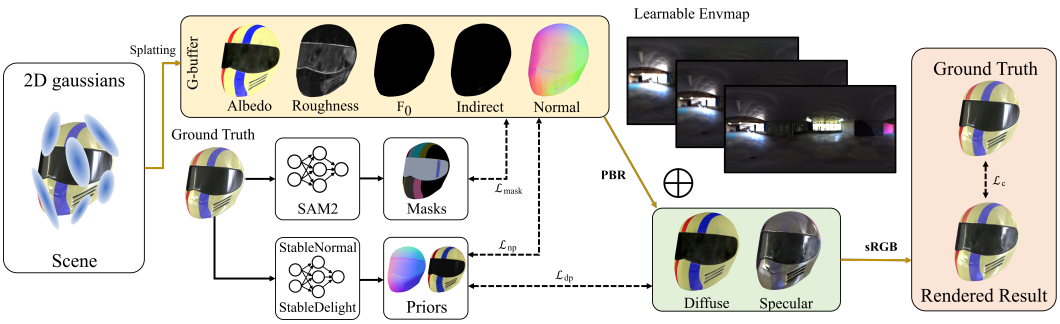


Figure 2: Overview of the MAGS framework.

and specular contributions to the rendered image. We retain the normal and diffuse-color priors used in SpecGloss-GS and use SAM2 segmentation masks to define regions for our mask-guided material consistency loss ($\mathcal{L}_{\text{mask}}$), which penalizes variation in rendered material attributes within each region.

Mask-Guided Material Regularization. Given a set of predefined segmentation masks $\{M_k\}_{k=1}^{K}$, where M_k denotes the set of pixels in the k-th mask, we first compute the mean material vector within each mask. Let $\mathbf{S}_i \in \mathbb{R}^5$ denote the rendered material vector at pixel i, consisting of RGB albedo, roughness, and metallic values. The mean material vector of the k-th mask is defined as

$$\boldsymbol{\mu}_k = \frac{1}{|M_k|} \sum_{i \in M_k} \mathbf{S}_i \tag{8}$$

Here, $|M_k|$ denotes the number of pixels contained in M_k. We then measure the material inconsistency within each mask by computing the average L_1 distance between the material vector of each pixel and the mean material vector of its corresponding mask:

$$\mathcal{L}_k = \frac{1}{|M_k|} \sum_{i \in M_k} \|\mathbf{S}_i - \boldsymbol{\mu}_k\| \tag{9}$$

Finally, the mask-guided material consistency loss is obtained by averaging these regional penalties across all masks:

$$\mathcal{L}_{\text{mask}} = \frac{1}{K} \sum_{k=1}^{K} \mathcal{L}_k \tag{10}$$

By minimizing $\mathcal{L}_{\text{mask}}$, material values within the same mask are encouraged to become closer to their corresponding mask-level mean, thereby promoting material consistency within each segmented region. Gradients propagate through the differentiable rasterization process to optimize the underlying Gaussian material attributes.

4 Experiments

In this section, we present the experimental setup and evaluation results of the proposed method. We evaluate its novel view synthesis quality and training efficiency on two benchmark datasets, followed by comparisons with representative Gaussian Splatting-based methods.

4.1 Experiment Settings

Datasets To evaluate novel view synthesis of reflective scenes, we conduct experiments on two benchmark datasets: Ref-Real and ShinyBlender (Verbin et al., 2024). Ref-Real consists of three captured real-world scenes, *Sedan*, *Garden Spheres*, and *Toycar*, featuring glossy surfaces and complex view-dependent reflections. ShinyBlender comprises six scenes, *Ball*, *Car*, *Coffee*, *Helmet*, *Teapot*, and *Toaster*, featuring glossy objects rendered in Blender under environment-map illumination. Together, these datasets allow us to assess the effectiveness of mask-guided material regularization in reproducing view-dependent reflective appearance in both real-world and synthetic settings.

Baseline We compare MAGS against five Gaussian Splatting-based approaches: 3DGS (Kerbl et al., 2023), 2DGS (Huang et al., 2024), EnvGS (Xie et al., 2025), MaterialRefGS (Zhang et al., 2026), and SpecGlossGS (Kouros et al., 2026). 3DGS and 2DGS serve as foundational baselines for evaluating improvements over standard radiance-based Gaussian representations, with 2DGS also providing the geometric backbone of MAGS. EnvGS and MaterialRefGS build upon 2DGS and address reflective novel view synthesis through explicit environment representation and material-aware rendering, respectively. SpecGlossGS provides the physically based rendering framework adopted by MAGS and serves as a direct baseline for assessing the additional benefit of our mask-guided material regularization.

Evaluation Metrics We evaluate the quality of synthesized novel views using three commonly adopted image quality metrics: Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index Measure (SSIM) (Wang et al., 2004), and Learned Perceptual Image Patch Similarity (LPIPS) (Zhang et al., 2018). We also measure the training time to evaluate the computational efficiency of each method, particularly the additional optimization cost introduced by material and reflection modeling.

Implementation details. We trained each Ref-Real scenes for 30k iterations and each ShinyBlender scenes for 50k iterations on an NVIDIA GeForce RTX 5090. Deferred shading is used throughout training. The indirect lighting stage begins at 10k iterations on Ref-Real scenes, 20k iterations on ShinyBlender scenes. After this point, we extract an object mesh every 3k iterations to enable visibility estimation via ray tracing. Every trainable attributes follow the settings of SpecGlossGS(Kouros et al., 2026). Loss weights are set to $\lambda = 0.2$, $\lambda_{\text{n}} = 0.03$, $\lambda_{\text{np}} = 0.01$, $\lambda_{\text{dp}} = 0.05$, $\lambda_{\text{light}} = 0.001$, and $\lambda_{\text{mask}} = 0.005$.

4.2 Comparisons

Comparisons on Ref-Real Dataset. We first evaluate our method on the real-world Ref-Real dataset (Verbin et al., 2024) and report the quantitative results in Table 1. On *Sedan*, MAGS achieves the best results among the compared methods across all three metrics, with a PSNR of 25.87 dB, an SSIM of 0.755, and an LPIPS of 0.236. On the other two scenes, MAGS maintains competitive performance, matching SpecGlossGS Kouros et al. (2026) on *Garden Spheres* and slightly improving PSNR and SSIM on *Toycar*. Across the dataset, MAGS achieves the second-highest average SSIM (0.679), while providing modest improvements in average PSNR and SSIM over SpecGlossGS Kouros et al. (2026) and maintaining comparable average LPIPS.

Comparisons on ShinyBlender. We also evaluate our method on the synthetic ShinyBlender dataset (Verbin et al., 2024) and report the quantitative results in Table 2. MAGS achieves the best or joint-best results among the compared methods for most scene–metric pairs. It also achieves the best average performance across all three metrics, with a PSNR of 35.51 dB, an SSIM of 0.978, and an LPIPS of 0.048. Compared with SpecGloss-GS, MAGS improves PSNR in five of the six scenes, with the largest gains on *Toaster* and *Teapot*, reaching 0.47 dB and 0.46 dB, respectively. On *Toaster*, this improvement is accompanied by a reduction in LPIPS from 0.078 to 0.071. The gains are scene-dependent, however: SpecGloss-GS performs better on *Coffee* and yields slightly lower LPIPS on *Car*, while EnvGS achieves higher PSNR on *Teapot*.

Table 1: Quantitative evaluation on the Ref-Real dataset (Verbin et al., 2024). The three best distinct scores are highlighted as 1st, 2nd, and 3rd, with tied scores sharing the same color.

	Garden Spheres	*Sedan*	*Toycar*	*Avg.*
		PSNR ↑		
2DGS	21.38	25.10	23.94	23.47
EnvGS	22.72	25.51	24.56	24.26
MaterialRefGS	23.26	25.32	24.80	24.46
SpecGloss-GS	22.52	25.76	24.15	24.14
Ours	22.52	25.87	24.17	24.19
		SSIM ↑		
2DGS	0.483	0.706	0.625	0.605
EnvGS	0.593	0.715	0.661	0.656
MaterialRefGS	0.643	0.709	0.690	0.681
SpecGloss-GS	0.619	0.753	0.663	0.678
Ours	0.619	0.755	0.664	0.679
		LPIPS ↓		
2DGS	0.443	0.364	0.315	0.374
EnvGS	0.205	0.320	0.234	0.253
MaterialRefGS	0.186	0.271	0.191	0.216
SpecGloss-GS	0.261	0.237	0.268	0.255
Ours	0.261	0.236	0.268	0.255

Table 2: Quantitative evaluation on the ShinyBlender dataset (Verbin et al., 2024). The three best distinct scores are highlighted as 1st, 2nd, and 3rd, with tied scores sharing the same color.

	Ball	*Car*	*Coffee*	*Helmet*	*Teapot*	*Toaster*	*Avg.*
			PSNR ↑				
2DGS	27.37	27.19	32.17	28.13	45.21	21.16	30.21
EnvGS	32.41	30.43	33.34	31.01	46.71	26.96	33.48
MaterialRefGS	27.63	29.21	34.44	27.61	44.13	21.52	30.76
SpecGloss-GS	37.60	32.06	34.59	34.42	45.95	27.44	35.34
Ours	37.76	32.07	34.20	34.68	46.41	27.91	35.51
			SSIM ↑				
2DGS	0.931	0.930	0.971	0.950	0.996	0.893	0.945
EnvGS	0.973	0.959	0.974	0.967	0.997	0.944	0.969
MaterialRefGS	0.947	0.951	0.976	0.950	0.996	0.889	0.952
SpecGloss-GS	0.988	0.973	0.976	0.980	0.997	0.949	0.977
Ours	0.988	0.973	0.975	0.981	0.997	0.952	0.978
			LPIPS ↓				
2DGS	0.175	0.050	0.079	0.083	0.008	0.131	0.088
EnvGS	0.121	0.038	0.078	0.055	0.006	0.078	0.063
MaterialRefGS	0.143	0.043	0.080	0.085	0.009	0.134	0.082
SpecGloss-GS	0.074	0.026	0.078	0.035	0.007	0.078	0.050
Ours	0.074	0.027	0.079	0.034	0.006	0.071	0.048

4.3 Qualitative Comparisons on Reflective Scenes

Figures 3 and 4 present qualitative comparisons on Ref-Real and ShinyBlender, respectively. MAGS achieves overall visual quality comparable to the competing methods, with localized refinements in



Figure 3: Qualitative comparison of MAGS with SpecGloss-GS, MaterialRefGS, and EnvGS on the *Sedan* (top) and *Garden Spheres* (bottom) scenes from the real-world Ref-Real dataset. The insets show enlarged reflective regions for comparing fine reflection details and their spatial alignment with the ground truth.

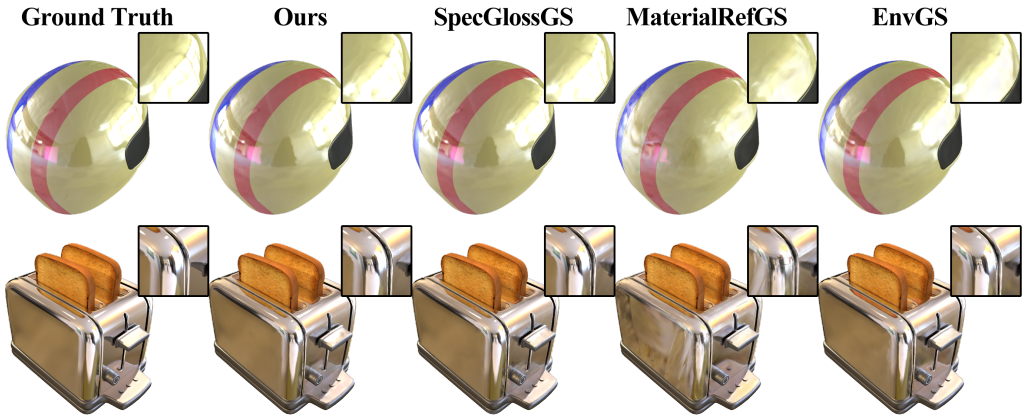


Figure 4: Qualitative comparison of MAGS with SpecGloss-GS, MaterialRefGS, and EnvGS on the *Helmet* (top) and *Toaster* (bottom) scenes from the synthetic ShinyBlender dataset. The insets show enlarged reflective regions for comparing fine reflection details and their spatial alignment with the ground truth.

reflective regions. In selected regions, reflected structures appear more coherent and continuous, with smoother spatial transitions in reflective appearance.

5 Conclusion

We propose MAGS, a framework for reflective novel view synthesis that addresses material inconsistencies among Gaussians within regions expected to share a common material. Building on physically based 2D Gaussian Splatting, we use segmentation masks as regional priors to regularize material attributes within each region, encouraging consistent material estimation for rendering view-dependent reflections.

Limitations. Our approach uses segmentation masks as proxies for regions with shared material properties. Inaccurate segmentation or heterogeneous materials within a region may therefore lead to over-regularization, suppressing genuine spatial variation in material attributes. Our evaluation is also limited to Ref-Real and Shiny Blender, and further experiments are needed to assess generalization to more diverse reflective scenes and material distributions.

AI USE STATEMENT

In this work, we used generative AI tools to assist with writing proofs, design or provide feedback on research methodology or experiments, implement methods, and assist with translation.

We did not use generative AI tools to generate synthetic datasets, develop theoretical models or conceptual frameworks, formulate mathematical claims, provide critical ingredients for proving mathematical claims, propose or refine hypotheses, clean or reformat datasets, support qualitative or thematic data analysis, or interpret results.

Additionally, we used generative AI tools to suggest experimental parameters, create or edit software code, summarize or analyze existing literature, edit the manuscript to improve readability, identify relevant literature, format references, suggest the manuscript structure, and propose a title or keywords. We did not use generative AI tools for the remaining tasks for which disclosure is recommended.

We have reviewed all AI-assisted work. We take responsibility for the final content of this work, including text, claims, and artifacts produced with the aid of generative AI.

REFERENCES

Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5460–5469. IEEE, 2022.

Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5491–5500. IEEE, 2022.

Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2d gaussian splatting for geometrically accurate radiance fields. In *ACM SIGGRAPH 2024 conference papers*, pp. 1–11, 2024.

Yingwenqi Jiang, Jiadong Tu, Yuan Liu, Xifeng Gao, Xiaoxiao Long, Wenping Wang, and Yuexin Ma. Gaussianshader: 3d gaussian splatting with shading functions for reflective surfaces. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5322–5332. IEEE, 2024.

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, George Drettakis, et al. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023.

Georgios Kouros, Minye Wu, Sushruth Nagesh, Xianling Zhang, and Tinne Tuytelaars. Unveiling the ambiguity in neural inverse rendering: A parameter compensation analysis. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pp. 2832–2841. IEEE, 2024.

Georgios Kouros, Minye Wu, and Tinne Tuytelaars. Spec-gloss surfels and normal-diffuse priors for relightable glossy objects. In *2026 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pp. 4417–4426. IEEE, 2026.

Shuichang Lai, Letian Huang, Jie Guo, Kai Cheng, Bowen Pan, Xiaoxiao Long, Jiangjing Lyu, Chengfei Lv, and Yanwen Guo. Glossygs: Inverse rendering of glossy objects with 3d gaussian splatting. *IEEE Transactions on Visualization and Computer Graphics*, 31(10):7478–7491, 2025.

Zhihao Liang, Qi Zhang, Ying Feng, Ying Shan, and Kui Jia. Gs-ir: 3d gaussian splatting for inverse rendering. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 21644–21653. IEEE, 2024.

Ben Mildenhall, Pratul P Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Transactions on Graphics (ToG)*, 38(4):1–14, 2019.

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. *Communications of the ACM*, 65(1):99–106, 2021.

Dor Verbin, Peter Hedman, Ben Mildenhall, Todd Zickler, Jonathan T Barron, and Pratul P Srinivasan. Ref-nerf: Structured view-dependent appearance for neural radiance fields. *IEEE transactions on pattern analysis and machine intelligence*, 47(11):9426–9437, 2024.

Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.

Tao Xie, Xi Chen, Zhen Xu, Yiman Xie, Yudong Jin, Yujun Shen, Sida Peng, Hujun Bao, and Xiaowei Zhou. Envgs: Modeling view-dependent appearance with environment gaussian. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5742–5751. IEEE, 2025.

Yuxuan Yao, Zixuan Zeng, Chun Gu, Xiatian Zhu, and Li Zhang. Reflective gaussian splatting. In *International Conference on Learning Representations*, volume 2025, pp. 68695–68711, 2025.

Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. Plenoctrees for real-time rendering of neural radiance fields. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 5732–5741. IEEE, 2021.

Zehao Yu, Anpei Chen, Binbin Huang, Torsten Sattler, and Andreas Geiger. Mip-splatting: Alias-free 3d gaussian splatting. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 19447–19456. IEEE, 2024.

Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *2018 IEEE/CVF conference on computer vision and pattern recognition*, pp. 586–595. IEEE, 2018.

Wenyuan Zhang, Jimin Tang, Weiqi Zhang, Yi Fang, Yu-Shen Liu, and Zhizhong Han. Materialrefgs: Reflective gaussian splatting with multi-view consistent material inference. *Advances in Neural Information Processing Systems*, 38:89896–89917, 2026.