Shading and Lighting Models in Computer Graphics

Goldi Soni	Ujjawal Singh	Saurabh Kumar Shriwastava
Assistant professor	B.Tech CSE	B.Tech CSE
Amity University Chhattisgarh, Raipur	Amity University Chhattisgarh, Raipur	Amity University Chhattisgarh, Raipur
gsoni@rpr.amity.edu	ujjwalsinghrathor70@gmail.com	saurabhshriwastava323214@gmail.com

Abstract — This review paper covers 30 important research papers published between 2021 and 2025 on the topic of shading and lighting models in computer graphics. These five years were one of the most exciting and fast-moving periods the field has ever seen, driven largely by the explosion of neural rendering techniques and the rapid advancement of real-time ray tracing hardware.

The papers reviewed here cover a wide range of topics. On the neural side, we look at Neural Radiance Fields (NeRF) and their descendants, which introduced a completely new way of representing and rendering 3D scenes using neural networks. We review how these methods were extended to handle material decomposition, relighting, and reflections. We also look at 3D Gaussian Splatting, a newer technique that achieves real-time rendering speeds while maintaining high visual quality.

On the real-time rendering side, we examine major advances in global illumination including ReSTIR GI, Lumen (Unreal Engine 5), and radiance caching systems. We also review neural denoising — the use of deep learning to clean up the noise produced by path tracing — which has become an essential part of every modern real-time renderer.

Each paper is summarized in plain, accessible language. For every paper, we describe what problem it solved, how the researchers approached the solution, and what impact the work had. No deep technical background is required to understand the summaries. Together, these 30 papers paint a picture of a field in rapid transformation, moving from manually engineered rendering algorithms toward hybrid systems that combine physics-based rendering with powerful machine learning.

Keywords- Neural Rendering, Global Illumination, Neural Radiance Fields (NeRF), Gaussian Splatting, Real-Time Rendering, Physically-Based Rendering (PBR), Deep Learning, Ray Tracing, Inverse Rendering, Shading and Lighting

1. Introduction

The years 2021 to 2025 will likely be remembered as a turning point in the history of computer graphics. For decades, the field progressed through careful mathematical modeling of how light interacts with surfaces. Researchers developed physically-based rendering (PBR) models, shadow mapping algorithms, global illumination techniques, and ever-more-efficient ray tracing methods. All of these advances were grounded in physics and expressed as equations that computer programs could evaluate.

Then, almost suddenly, a new approach arrived: neural rendering. Starting with NeRF (Neural Radiance Fields) in 2020 and exploding into hundreds of follow-up papers in 2021 and beyond, researchers discovered that neural networks — particularly deep learning models — could learn to represent and render 3D scenes from photographs with staggering quality. These neural methods did not replace physics; instead they learned to approximate physical light behavior from data, often producing results that looked better than anything achievable with hand-designed equations.

At the same time, the real-time rendering world was transformed by two developments: the arrival of dedicated ray tracing hardware in consumer graphics cards (GPUs), and the introduction of techniques like ReSTIR that made sampling complex lighting dramatically more efficient. These changes made it possible to bring effects like global illumination and soft shadows — previously only possible in slow offline renderers used for movies — into real-time video games.

This review paper covers both of these threads. The 30 papers selected represent a mixture of neural rendering breakthroughs (NeRF variants, Gaussian splatting, neural materials), real-time rendering advances (ReSTIR GI, Lumen, radiance caching), and work that bridges the two worlds (neural denoising for path tracing, neural inverse rendering, AI-guided lighting). Together they give a comprehensive picture of where the field of shading and lighting stood between 2021 and 2025.

Each paper summary follows the same format: title, authors, year, objective (what problem it solved), methodology (how it solved it), and conclusions (what it achieved and why it matters). The language is kept as simple as possible through

2. Literature review

This paper focuses on solving the problem of relighting objects from images using neural rendering techniques. The main objective is to separate lighting, material properties, and visibility for realistic image synthesis. The methodology extends Neural Radiance Fields by incorporating reflectance modeling and visibility prediction into the neural network. The system learns surface roughness, shadow regions, and indirect lighting using training images. It applies physics-based rendering equations along with neural approximations. The model estimates how light interacts with surfaces under different conditions. Experimental results show high-quality relighting even under unseen lighting environments. The system successfully distinguishes between shadows and intrinsic object properties. This improves realism and flexibility in rendering pipelines. The contribution lies in integrating neural networks with physically-based rendering concepts. Itenables advanced editing such as changing lighting without altering geometry. The study also lays a foundation for neural inverse rendering. It has strong applications in gaming, film production, and virtual reality. Overall, it represents a significant step toward intelligent lighting models.[1]

This paper addresses the challenge of decomposing object appearance into material and lighting components. The objective is to recover surface reflectance properties from multiple images. The

methodology uses neural volumetric representations similar to NeRF. It models albedo, specular reflection, and surface normals within a neural framework. Lighting is represented using spherical Gaussian functions for efficiency. The system is trained end-to-end using multi-view image datasets. It simultaneously learns geometry, materials, and lighting conditions. Results show accurate decomposition and high-quality relighting capabilities. The method allows objects to be placed in new lighting environments realistically. It also enables exporting results into standard 3D formats for real-time applications. The contribution is significant for material editing and scene reconstruction. It reduces dependence on manual modeling techniques. The study improves automation in 3D content creation pipelines. It is useful in gaming, AR, and digital production. Overall, it enhances realism and flexibility in rendering workflows.[2]

This paper introduces a novel way of representing 3D scenes using neural networks. The objective is to generate photorealistic images from 2D photographs. The methodology encodes 3D coordinates and viewing directions into a neural network. The network predicts color and density values for each point in space. Rendering is performed using volumetric integration along rays. The model is trained using image reconstruction loss from multiple views. Results demonstrate highly realistic novel view synthesis. The method captures fine geometric and lighting details. The contribution revolutionized scene representation in computer graphics. It replaced traditional mesh-based methods with implicit neural representations. However, initial models were computationally expensive. Despite limitations, it became the foundation for many future works. It enabled advancements in view synthesis, relighting, and scene understanding. The study has widespread applications in VR, AR, and film production. Overall, it is one of the most influential papers in modern graphics.[3]

This paper improves the rendering of reflective and shiny surfaces in neural models. The objective is to accurately represent view-dependent effects such as reflections. The methodology introduces reflection-based parameterization instead of traditional view-direction input. It separates diffuse and specular components for better control. The neural network predicts surface normals to compute reflection directions. This improves the physical correctness of reflections. Results show sharper highlights and reduced artifacts. The method successfully handles complex reflective materials. It enhances realism in neural rendering systems. The contribution lies in improving specular rendering accuracy. It also improves interpretability of neural scene representations. The approach is widely adopted in later works. It supports applications in product visualization and virtual environments. Overall, it strengthens the capability of neural rendering for realistic materials.[4]

This paper focuses on accelerating neural rendering techniques. The objective is to reduce training and rendering time significantly. The methodology introduces multiresolution hash encoding for fast data representation. Instead of heavy neural networks, it uses compact feature lookups. This drastically reduces computational cost. Training time is reduced from hours to seconds. Rendering is achieved in real-time with high quality. Results show comparable performance to NeRF with much higher speed. The contribution makes neural rendering practical for real-world applications. It enables interactive systems such as games and simulations. The method is widely adopted in industry tools. It bridges the gap between research and production. The study also supports scalability for large scenes. Overall, it is a breakthrough in efficient neural rendering.[5]

This paper addresses the problem of reconstructing 3D assets from images. The objective is to extract geometry, materials, and lighting simultaneously. The methodology uses differentiable rendering for optimization. Geometry is represented using signed distance fields for accuracy. Materials follow physically-based rendering parameters. Lighting is modeled using environment maps. The system is trained to minimize differences between rendered and real images. Results show accurate reconstruction of complex objects. The output is compatible with standard 3D pipelines. The contribution enables automated asset creation. It reduces manual effort in modeling and texturing. The study is useful for game development and visual effects. It integrates AI with traditional rendering techniques. Overall, it advances inverse rendering capabilities.[6]

This paper focuses on real-time global illumination. The objective is to simulate indirect lighting efficiently. The methodology uses reservoir-based sampling techniques. It reuses samples across spatial and temporal domains. This reduces noise while maintaining accuracy. The system shares light paths between pixels. Results show significant improvement in rendering quality. It works effectively with limited computational resources. The contribution enables real-time path tracing. It is widely used in modern rendering engines. The study improves lighting realism in games and simulations. It balances quality and performance effectively. The technique is scalable for complex scenes. Overall, it represents a major advancement in global illumination.[7]

This paper introduces a real-time global illumination system for Unreal Engine 5. The objective is to provide dynamic lighting without precomputation. The methodology combines ray tracing, radiance caching, and screen-space techniques. It supports both hardware and software ray tracing. The system dynamically updates lighting based on scene changes. Results show high-quality lighting in real-time applications. It handles reflections, shadows, and indirect lighting effectively. The contribution is widely adopted in game development. It eliminates the need for baked lighting. The study improves workflow efficiency for developers. It supports large-scale dynamic environments. Overall, it sets a new standard for real-time rendering systems.[8]

This paper introduces a new representation for neural rendering. The objective is to achieve real-time rendering with high quality. The methodology represents scenes using 3D Gaussian primitives. Rendering is performed using GPU-based splatting techniques. The system learns Gaussian parameters from images. Results show real-time performance at high resolution. The quality is comparable or better than NeRF. The contribution provides a faster alternative to neural fields. It supports interactive applications. The method is widely adopted in research and industry. It improves scalability and efficiency. The study opens new directions in scene representation. Overall, it is a major advancement in real-time rendering.[9]

This paper extends Gaussian splatting for inverse rendering. The objective is to separate geometry, materials, and lighting. The methodology adds physically-based material parameters to Gaussian primitives. It estimates surface normals and environment lighting. The system is trained using multi-view images. Results allow relighting and material editing. The contribution improves flexibility in scene manipulation. It enables realistic editing of rendered scenes. The study combines speed with physical accuracy. It supports interactive applications. The method is efficient and scalable. Overall, it enhances Gaussian-based rendering systems.[10]

This paper focuses on efficient rendering of complex material appearances in real-time systems. The objective is to replace computationally expensive shading models with compact neural representations. The methodology uses small neural networks to approximate bidirectional reflectance distribution functions (BRDFs). These networks are trained offline on high-quality material datasets. The system transforms input directions into learned shading frames for better accuracy. It also incorporates importance sampling techniques for efficient rendering. Results show that neural models can reproduce complex materials such as fabrics, layered coatings, and glass. The performance is significantly faster compared to traditional shading methods. The contribution lies in compressing complex material behavior into lightweight neural networks. This reduces memory usage and computation time. It enables real-time rendering in games and simulations. The approach is compatible with modern GPU pipelines. Overall, it advances real-time material rendering using AI techniques.[11]

This paper provides a theoretical foundation for sampling techniques used in rendering. The objective is to formalize the ReSTIR method mathematically. The methodology introduces Generalized Resampled Importance Sampling (GRIS). It explains how reservoir sampling can be applied correctly across pixels and frames. The framework ensures unbiased rendering under defined conditions. It also extends sampling to multi-bounce lighting scenarios. Results validate the correctness and efficiency of the approach. The contribution strengthens the reliability of real-time rendering systems. It allows developers to safely implement advanced sampling methods. The study improves the theoretical understanding of rendering algorithms. It also supports future research in sampling optimization. The framework is widely used in modern graphics pipelines. Overall, it bridges the gap between theory and practice.[12]

This paper addresses efficient importance sampling for neural materials. The objective is to reduce noise in rendering neural BRDFs. The methodology uses neural networks combined with normalizing flows. These networks learn probability distributions of light directions. The system samples directions where light contribution is highest. It is trained on real-world material datasets. Results show significant reduction in rendering noise. The method improves efficiency of path tracing systems. The contribution enables neural materials to be used in practical rendering pipelines. It eliminates the need for manual sampling strategies. The study enhances integration of AI in rendering workflows. It is useful for photorealistic rendering applications. Overall, it improves both quality and performance.[13]

This paper improves traditional precomputed radiance transfer (PRT) methods. The objective is to reduce memory usage while increasing flexibility. The methodology replaces large precomputed tables with neural networks. The network predicts light transport for each surface point. It takes position and light direction as inputs. Results show improved memory efficiency and higher-quality lighting. The system supports dynamic lighting conditions better than classical PRT. The contribution modernizes precomputed rendering using AI. It allows real-time rendering of complex lighting effects. The study integrates neural methods with classical techniques. It is useful for games and simulations. Overall, it enhances efficiency and scalability of PRT methods.[14]

This paper focuses on relighting outdoor scenes using neural rendering. The objective is to simulate realistic lighting including sunlight and shadows. The methodology uses neural networks trained on images captured at different times of day. It models albedo, shadows, and lighting separately. The system estimates sun position and sky illumination. Results show accurate relighting with realistic shadow behavior. The method handles large-scale outdoor environments effectively. The contribution enables time-of-day editing in scenes. It is useful for visual effects and architectural visualization. The study improves realism in outdoor rendering. It reduces dependence on manual lighting setup. Overall, it advances neural relighting techniques.[15]

This paper proposes a scalable global illumination solution. The objective is to support real-time GI across different hardware. The methodology uses a two-level radiance caching system. It combines world-space probes with screen-space refinement. Rays update lighting information over time. Results show stable and efficient GI performance. The system adapts to dynamic scenes. The contribution improves scalability of GI techniques. It supports large environments in games. The study focuses on practical implementation. It balances quality and performance effectively. Overall, it enhances real-time lighting systems.[16]

This paper introduces reinforcement learning for rendering optimization. The objective is to improve sampling efficiency in path tracing. The methodology trains neural networks to allocate rays adaptively. It uses reinforcement learning to guide sampling decisions. A denoiser is also integrated into the system. Results show reduced rendering time with maintained quality. The contribution combines sampling and denoising into one framework. It improves efficiency of rendering pipelines. The study demonstrates AI-driven optimization. It is useful for real-time applications. Overall, it enhances intelligent rendering systems.[17]

This paper focuses on improving denoising quality. The objective is to preserve fine details while removing noise. The methodology uses a multi-scale pyramid structure. Each level processes different image details. Neural networks denoise at each scale separately. Results show sharper images compared to traditional methods. The contribution improves detail preservation. It enhances visual quality in path tracing. The method is efficient and scalable. It is easy to integrate into pipelines. Overall, it advances neural denoising techniques.[18]

This paper extends Gaussian splatting for relighting. The objective is to support dynamic lighting changes. The methodology adds material properties to Gaussian primitives. It uses ray tracing for shadows and reflections. Results show realistic relighting in real-time. The contribution combines speed and flexibility. It supports editing of lighting conditions. The study improves Gaussian rendering capabilities. It is useful for AR and VR applications. Overall, it enhances real-time relighting systems.[19]

This paper improves reflection rendering in Gaussian splatting. The objective is to handle shiny surfaces accurately. The methodology uses deferred shading techniques. It computes reflections using screen-space data. Results show improved reflection quality. The contribution solves limitations of Gaussian methods. It enhances realism in rendering. The method is efficient and scalable. Overall, it improves visual fidelity.[20]

This paper focuses on lighting control in AI-generated images. The objective is to allow precise lighting manipulation. The methodology uses diffusion models with lighting inputs. It trains on datasets with known lighting conditions. Results show controllable image generation. The contribution integrates AI with lighting design. It is useful for artists and designers. Overall, it improves generative rendering systems.[21]

This paper predicts indoor lighting from a single image. The objective is to estimate spatial lighting variations. The methodology uses graph neural networks. It models relationships between scene regions. Results show accurate lighting prediction. The contribution supports augmented reality applications. It improves realism in virtual object placement. Overall, it enhances lighting estimation.[22]

This paper uses GANs to approximate global illumination. The objective is to reduce computational cost. The methodology splits lighting into components. Each component is learned using neural networks. Results show fast rendering with acceptable quality. The contribution replaces expensive GI calculations. It is useful for real-time applications. Overall, it improves efficiency.[23]

This paper models light using wave optics. The objective is to simulate fine surface effects. The methodology uses electromagnetic equations. It captures diffraction and interference. Results show highly accurate reflectance. The contribution improves material realism. It is useful for research and validation. Overall, it enhances physical accuracy.[24]

This paper compares different GI techniques. The objective is to evaluate performance and quality. The methodology benchmarks multiple rendering systems. Results show trade-offs between speed and accuracy. The contribution guides developers in choosing methods. It supports practical decision-making. Overall, it improves understanding of GI systems.[25]

This paper improves inverse rendering speed. The objective is faster material and lighting estimation. The methodology uses tensor-based representations. Results show improved efficiency over NeRF. The contribution accelerates inverse rendering. It is useful for real-time applications. Overall, it enhances performance.[26]

This paper focuses on rendering in VR environments. The objective is to reduce computational cost. The methodology uses foveated rendering. Results show performance improvement. The contribution enhances VR rendering efficiency. Overall, it improves user experience.[27]

This paper improves reflection rendering in NeRF. The objective is consistent reflections. The methodology uses secondary neural representations. Results show stable reflections. The contribution enhances realism. Overall, it improves neural rendering.[28]

This paper adds global illumination to Gaussian scenes. The objective is dynamic lighting. The methodology uses radiance caching and ray tracing. Results show realistic GI. The contribution enhances Gaussian rendering. Overall, it improves realism.[29]

This paper introduces transformer-based rendering. The objective is neural global illumination. The methodology uses attention mechanisms. Results show fast rendering. The contribution bridges neural and traditional rendering. Overall, it represents future direction.[30]

3. Objectives

The primary objective of this review paper is to analyse and synthesize existing research on shading and lighting models in computer graphics, with a focus on recent advancements in neural rendering, real-time global illumination, and AI-driven techniques. Based on the reviewed literature, the study aims to achieve the following specific objectives:

- To examine various shading and lighting techniques, including physically-based rendering, ray tracing, neural rendering, and hybrid approaches used in modern computer graphics.
- To understand how machine learning and deep learning methods improve lighting simulation in terms of realism, computational efficiency, and automation.
- To analyse different approaches such as global illumination, neural radiance fields, Gaussian splatting, inverse rendering, and neural denoising techniques.
- To compare traditional rendering methods with modern AI-driven techniques and evaluate their advantages, limitations, and performance trade-offs.
- To study real-time rendering solutions used in gaming and virtual environments, focusing on efficiency, scalability, and visual quality.

- To evaluate the role of sampling techniques, radiance caching, and denoising methods in improving rendering performance and reducing noise.

- To identify challenges and limitations in current shading and lighting models, including computational cost, scalability issues, and handling of complex materials.

- To explore applications of advanced lighting models in areas such as gaming, virtual reality, augmented reality, film production, and digital content creation.

- To highlight emerging trends such as neural rendering, transformer-based rendering, and AI-guided lighting systems.

• To propose future research directions for developing efficient, real-time, and intelligent rendering systems.

4. Comparison Of Past Published Research Paper

S.No	Title of Paper	Year	Proper Objective	Methodology	Conclusion / Result
1	NeRF: Representing Scenes as Neural Radiance Fields	2021	Generate realistic 3D scenes from 2D images	Neural network (MLP) + volume rendering	Produces highly realistic novel views and revolutionized neural rendering
2	Instant Neural Graphics Primitives (Instant NGP)	2022	Speed up neural rendering for real-time use	Multiresolution hash encoding + GPU acceleration	Reduces training time from hours to seconds
3	ReSTIR GI: Real-Time Path Tracing	2021	Improve real-time global illumination quality	Reservoir-based sampling + spatiotemporal reuse	Enables high-quality indirect lighting in real-time
4	Lumen: Real-Time Global Illumination (UE5)	2022	Achieve fully dynamic GI in games	Hybrid ray tracing + radiance caching	Makes real-time global illumination practical in game engines
5	3D Gaussian Splatting for Real-Time Rendering	2023	Enable fast and high-quality scene rendering	Gaussian-based scene representation + GPU rasterization	Achieves real-time rendering with high visual quality

5. Conclusion

This review paper highlights the significant advancements in shading and lighting models in computer graphics, emphasizing the transition from traditional physically-based rendering techniques to modern AI-driven and neural rendering approaches. Methods such as Neural Radiance Fields, 3D Gaussian Splatting, and real-time global illumination have greatly improved visual realism, computational efficiency, and rendering speed. These technologies enable accurate simulation of light interactions, material properties, and dynamic environments, making them highly suitable for applications in gaming, virtual reality, and film production.

Overall, the integration of machine learning with traditional rendering techniques has transformed the field, making rendering systems more efficient, flexible, and capable of producing high-quality results in real time.

6.Future Scope

The future of shading and lighting models in computer graphics lies in the deeper integration of artificial intelligence with traditional rendering techniques. Future research is expected to focus on developing lightweight and efficient neural models that can perform high-quality rendering in real time with lower computational cost. Improvements in hardware acceleration and optimization algorithms will further enhance performance and scalability for large and dynamic scenes.

Emerging techniques such as transformer-based rendering, neural global illumination, and AI-guided lighting control have the potential to revolutionize how lighting is designed and simulated. Additionally, hybrid rendering approaches that combine physically-based methods with machine learning will provide better accuracy and flexibility. Future work will also focus on improving material representation, handling complex lighting conditions, and enabling real-time editing and interaction.

Overall, advancements in this field will lead to more realistic, efficient, and intelligent rendering systems, with wide applications in gaming, virtual reality, augmented reality, and digital content creation.

References

[1] Srinivasan, P. P., Deng, B., Zhang, X., Tancik, M., Mildenhall, B., & Barron, J. T. (2021). *NeRV: Neural reflectance and visibility fields for relighting and view synthesis*. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).

[2] Boss, M., Braun, R., Jampani, V., Barron, J. T., Liu, C., & Lensch, H. (2021). *NeRD: Neural reflectance decomposition from image collections*. Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV).

[3] Mildenhall, B., Srinivasan, P. P., Tancik, M., Barron, J. T., Ramamoorthi, R., & Ng, R. (2021). *NeRF: Representing scenes as neural radiance fields for view synthesis*. Communications of the ACM, 65(1), 99–106.

[4] Verbin, D., Hedman, P., Mildenhall, B., Zickler, T., Barron, J. T., & Srinivasan, P. P. (2022). *Ref-NeRF: Structured view-dependent appearance for neural radiance fields*. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).

[5] Müller, T., Evans, A., Schied, C., & Keller, A. (2022). *Instant neural graphics primitives with a multiresolution hash encoding*. ACM Transactions on Graphics, 41(4), 1–15.

[6] Munkberg, J., Hasselgren, J., Shen, T., Gao, J., Chen, W., Evans, A., Müller, T., & Fidler, S. (2022). *Extracting triangular 3D models, materials, and lighting from images*. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).

[7] Ouyang, Y., Liu, S., Kettunen, M., Pharr, M., & Pantaleoni, J. (2021). *ReSTIR GI: Path resampling for real-time path tracing*. ACM High Performance Graphics (HPG).

[8] Wright, D., Narkowicz, K., & Kelly, P. (2022). *Lumen: Real-time global illumination in Unreal Engine 5*. ACM SIGGRAPH Advances in Real-Time Rendering Course.

[9] Kerbl, B., Kopanas, G., Leimkühler, T., & Drettakis, G. (2023). *3D Gaussian splatting for real-time radiance field rendering*. ACM Transactions on Graphics, 42(4).

[10] Liang, Z., Zhang, Q., Feng, Y., Shan, Y., & Jia, K. (2024). *GS-IR: 3D Gaussian splatting for inverse rendering*. Proceedings of CVPR.

[11] Zeltner, T., Rousselle, F., Weidlich, A., Clarberg, P., Novák, J., Bitterli, B., Evans, A., Davidovič, T., Kallweit, S., & Lefohn, A. (2024). *Real-time neural appearance models*. ACM Transactions on Graphics.

[12] Lin, D., Kettunen, M., Bitterli, B., Pantaleoni, J., Yuksel, C., & Wyman, C. (2022). *Generalized resampled importance sampling: Foundations of ReSTIR*. ACM Transactions on Graphics.

[13] Xu, B., Wu, L., Hasan, M., Luan, F., Georgiev, I., Xu, Z., & Ramamoorthi, R. (2023). *NeuSample: Importance sampling for neural materials*. ACM SIGGRAPH.

[14] Rainer, G., Bousseau, A., Ritschel, T., & Drettakis, G. (2022). *Neural precomputed radiance transfer*. Computer Graphics Forum.

[15] Rudnev, V., Elgharib, M., Smith, W., Liu, L., Golyanik, V., & Theobalt, C. (2022). *NeRF-OSR: NeRF for outdoor scene relighting*. European Conference on Computer Vision (ECCV).

[16] Boissé, G., Meunier, S., de Dinechin, H., Bartels, P., Lauritzen, A., & Lorach, T. (2022). *GI-1.0: A fast scalable two-level radiance caching scheme for real-time global illumination*. AMD Technical Report.

[17] Scardigli, A., Cavigelli, L., & Müller, L. K. (2023). *RL-based stateful neural adaptive sampling and denoising for real-time path tracing*. Advances in Neural Information Processing Systems (NeurIPS).

[18] Balint, M., Wolski, K., Myszkowski, K., Seidel, H. P., & Mantiuk, R. (2023). *Neural partitioning pyramids for denoising Monte Carlo renderings*. ACM SIGGRAPH.

[19] Gao, J., Gu, C., Lin, Y., Zhu, H., Cao, X., Zhang, L., & Yao, Y. (2024). *Relightable 3D Gaussian: Real-time point cloud relighting with BRDF decomposition and ray tracing*. ECCV.

[20] Ye, K., Hou, Q., & Zhou, K. (2024). *3D Gaussian splatting with deferred reflection*. ACM SIGGRAPH.

[21] Zeng, C., Dong, Y., Peers, P., Kong, Y., Wu, H., & Tong, X. (2024). *DiLightNet: Fine-grained lighting control for diffusion-based image generation*. ACM SIGGRAPH.

[22] Bai, J., Guo, J., Wang, C., Li, M., Chen, Z., Pan, J., Guo, Y., & Yan, L. Q. (2023). *Deep graph learning for spatially-varying indoor lighting prediction*. Science China Information Sciences.

[23] Sousa, E. V., & Fernandes, L. A. F. (2023). *Approximating global illumination using deep learning with generative adversarial networks*. Pattern Recognition Letters.

[24] Li, P., Luan, F., Zhao, S., Gkioulekas, I., & Zhang, C. (2023). *A full-wave reference simulator for computing surface reflectance*. ACM SIGGRAPH.

[25] Lambru, C., Morar, A., Moldoveanu, F., Asavei, V., & Moldoveanu, A. (2021). *Comparative analysis of real-time global illumination techniques in current game engines*. IEEE Access.

[26] Jin, H., Liu, I., Xu, P., Zhang, X., Han, S., Bi, S., Zhou, X., Xu, Z., & Su, H. (2023). *TensoIR: Tensorial inverse rendering*. CVPR.

[27] Peres, V., Brito, C. J. S., Gois, J. P., & Boechat, P. (2023). *Non-homogeneous denoising for virtual reality in real-time path tracing rendering*. Computers & Graphics.

[28] Verbin, D., Srinivasan, P. P., Hedman, P., Mildenhall, B., Poole, B., Barron, J. T., & Szeliski, R. (2024). *NeRF-Casting: Improved view-dependent appearance with consistent reflections*. SIGGRAPH Asia.

[29] Hu, C., Gai, M., Wang, G., & Li, S. (2025). *Real-time global illumination for dynamic 3D Gaussian scenes*. arXiv.

[30] Zeng, C., Guo, B., Dong, Y., et al. (2025). *RenderFormer: Transformer-based neural rendering of triangle meshes with global illumination*. arXiv.