Published by :
https://www.ijert.org/
An International Peer-Reviewed Journal

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

Neural Rendering in Computer Graphics

Goldi Soni
Assistant Professor
Amity Univesity, Chhattisgarh

Shilpa Vishwakarma
B.Tech(CSE)
Amity Univesity, Chhattisgarh

Hina Sahu
B.Tech(CSE)
Amity Univesity, Chhattisgarh

Abstract - Neural rendering has emerged as a transformative paradigm that bridges traditional computer graphics with deep learning, enabling photorealistic image synthesis and novel view generation from limited data. This review paper systematically surveys 30 research works spanning 2020 to 2026, covering foundational techniques, accelerator-driven optimisations, and domain-specific applications of neural rendering. Key methodologies examined include Neural Radiance Fields (NeRF), flow matching, deferred rendering, transformer-based pipelines, and light-field representations. Hardware acceleration strategies—such as domain-specific coprocessors, unified accelerator architectures, and mobile-optimised redundancy elimination—are critically analysed. Domain applications explored range from autonomous driving and medical simulation to underwater robotics, human motion reconstruction, and game character auto-creation. Despite remarkable progress, challenges persist in achieving real-time inference, scalable generalisation, and memory efficiency. This review synthesises current findings and identifies future directions including hardware-software co-design, lightweight model compression, and integration with physical simulation engines.

Keywords: Neural Rendering, Neural Radiance Fields (NeRF), Real-Time Rendering, Deep Learning, Hardware Acceleration

1. INTRODUCTION

Neural rendering represents the integration of machine learning—particularly deep neural networks—with conventional computer graphics pipelines. Traditional rendering techniques rely on physically based simulation of light transport and hand-crafted material models, which, while accurate, are computationally expensive and struggle with real-time constraints. Neural rendering overcomes these limitations by learning compact, data-driven scene representations that can synthesise photorealistic images efficiently.

The seminal introduction of Neural Radiance Fields (NeRF) in 2020 catalysed a rapid expansion of research in this area. NeRF-based methods model a scene as a continuous volumetric function mapping 3D positions and viewing directions to colour and density, enabling novel view synthesis with remarkable fidelity. Subsequent works have extended this foundation to support real-time performance, dynamic scenes, generalisable representations, and hardware acceleration.

This review paper examines 30 significant publications in neural rendering published between 2020 and 2026. The scope includes algorithmic innovations, hardware acceleration frameworks, and application-specific adaptations. By contextualising each work within the broader trajectory of the field, this paper aims to provide a comprehensive reference for researchers and practitioners working at the intersection of computer graphics and deep learning.

1.1 Neural Rendering Pipeline — Conceptual Overview

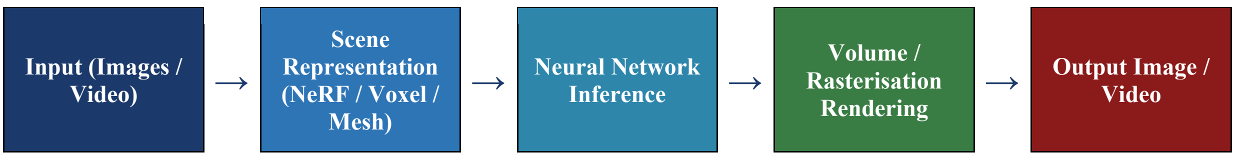

Figure 1 illustrates the general workflow of a neural rendering system.

1.2 Taxonomy of Neural Rendering Approaches

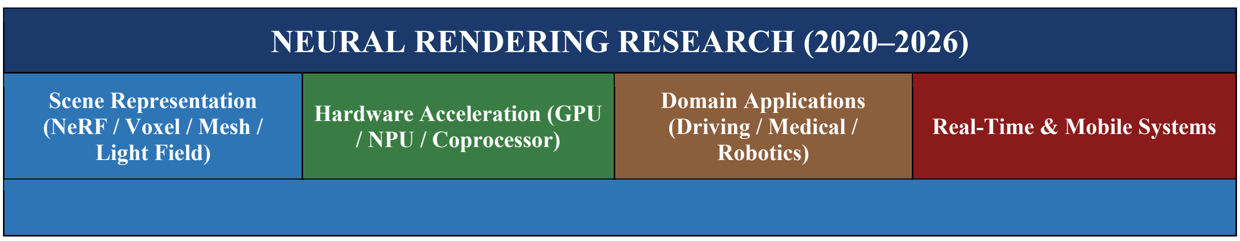

Figure 2 presents a flowchart classifying the major categories of neural rendering research reviewed in this paper.

2. LITERATURE REVIEW

RenderFlow, proposed by Shenghao Zhang, Runtao Liu, Christopher Schroers, and Yang Zhang (2026), introduces a fast and deterministic neural rendering framework designed to overcome the latency and randomness of diffusion-based rendering methods. The approach reformulates neural rendering as a single-step flow matching problem, directly mapping albedo images to fully shaded outputs using G-buffer information such as depth, normals, and material properties. Unlike traditional diffusion models, RenderFlow avoids iterative noise removal, enabling significantly faster rendering suitable for real-time and interactive applications. The framework leverages a pretrained video Diffusion Transformer backbone and applies a bridge matching strategy in latent space for efficient training. To improve physical consistency and temporal stability, the method incorporates optional sparse keyframe guidance based on high-quality path-traced reference frames. Additionally, an adapter-based inverse rendering module allows intrinsic scene decomposition without retraining the full model. However, the approach relies heavily on large-scale synthetic datasets generated using Unreal Engine, which may limit real-world generalization. Training also requires significant computational resources and high-performance GPUs. Overall, RenderFlow demonstrates that single-step flow matching can provide an efficient alternative to diffusion-based neural rendering while maintaining photorealistic quality.[1]

Adaptive Neural Rendering for Dynamic Advanced Life Support Team Training (2026), proposed by Basel Salem Alghamdy and colleagues, presents an adaptive neural rendering framework aimed at improving Advanced Life Support (ALS) team training through real-time simulation scenario adaptation. The approach focuses on enhancing the realism and responsiveness of medical training environments by dynamically adjusting scenarios based on trainee actions and decision-making. The authors introduce a continuous volumetric neural rendering framework capable of modeling complex 3D medical environments while enabling real-time updates to patient conditions, environmental factors, and procedural difficulty. By continuously updating volumetric representations, the system can simulate realistic physiological and environmental changes during training sessions. This adaptive mechanism allows training simulations to respond interactively to team performance, improving learning effectiveness. However, the framework requires significant computational resources to support real-time rendering and continuous updates, resulting in high hardware dependency. In addition, implementation complexity and the need for extensive validation with medical professionals may limit large-scale deployment. Despite these limitations, the study demonstrates that adaptive neural rendering with continuous volumetric representation can significantly improve the flexibility and realism of ALS training simulations. Overall, the framework shows strong potential for advancing real-time medical training systems and enhancing emergency response preparedness. [2]

Real-Time Robot Self-Modeling via Direct Sensorimotor Decoding: Overcoming the Neural Rendering Latency Bottleneck (2026), proposed by Muhammad Zeeshan Asghar, presents a framework for enabling real-time robot self-modeling while addressing latency issues in neural rendering pipelines. The primary objective of the work is to allow robots to continuously learn and update their internal body models during operation. The author introduces a direct sensorimotor decoding framework that bypasses explicit visual or 3D reconstruction stages. Instead, neural networks directly map sensor inputs and motor commands to self-model parameters. This approach significantly reduces computational overhead and enables faster model updates. As a result, robots can adapt their internal representations more efficiently in real-time environments. However, the method sacrifices visual interpretability since explicit rendering is avoided. The framework also depends heavily on the quality and diversity of sensorimotor data. In addition, its generalization to complex robot morphologies and dynamic environments may require further training and validation. Overall, the study demonstrates that direct sensorimotor decoding can effectively overcome neural rendering latency and improve the adaptability and responsiveness of robotic systems.[3]

Neural Scene Baking for Permutation Invariant Transparency Rendering with Real-time Global Illumination (2026), proposed by Ziyang Zhang and Edgar Simo-Serra, presents a neural rendering framework designed to enable real-time rendering of transparent objects with global illumination. The primary goal of the work is to eliminate the order dependence commonly found in traditional transparency rendering techniques, which often causes visual artifacts and increased computational cost. The authors introduce a neural scene baking framework that precomputes global illumination effects using neural networks and encodes complex light transport phenomena such as reflection and refraction through transparent materials. A permutation-invariant neural architecture ensures consistent rendering results regardless of the processing order of transparent surfaces. This learned scene representation allows efficient real-time rendering during inference while maintaining high visual realism. However, the approach requires offline preprocessing and training, which can be time-consuming. Its performance also depends on the quality and diversity of the training dataset,

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

potentially limiting generalization to unseen scenes. Additionally, large-scale scenes with many transparent objects may increase memory requirements. Overall, the method provides an effective solution for real-time transparency rendering with global illumination, making it valuable for interactive graphics, virtual reality, and gaming applications.[4]

An Overview of Neural Rendering Accelerators: Challenges, Trends, and Future Directions (2025), authored by Junha Ryu and Hoi-Jun Yoo, presents a comprehensive survey of hardware accelerators designed for neural rendering workloads. The paper aims to bridge the gap between neural rendering algorithms and efficient hardware implementations. The authors review several hardware platforms, including Graphics Processing Unit, Neural Processing Unit, and specialized neural rendering accelerators, evaluating them based on performance, energy efficiency, and scalability. Key challenges such as memory bandwidth limitations, latency, and excessive data movement are analyzed in detail. The survey also discusses optimization techniques including model compression, pipeline parallelism, and hardware–software co-design to improve real-time performance. As a survey-based study, the paper does not propose a new accelerator architecture and relies on results from previously published works. Moreover, rapid developments in hardware technologies may affect the long-term applicability of some discussed approaches. Overall, the study highlights the importance of specialized accelerator designs for neural rendering tasks and outlines important trends and future research directions for next-generation real-time graphics systems.[5]

Neural Rendering Coprocessor With Optimized Ray Representation and Marching (2025), proposed by Zhechen Yuan and colleagues, presents a dedicated hardware accelerator designed to improve the efficiency of neural rendering workloads. The main objective of the study is to reduce computational overhead and memory access costs associated with ray-based rendering operations. The authors introduce a hardware–software co-designed neural rendering coprocessor that utilizes a compact ray representation to efficiently encode ray information. In addition, an optimized ray marching pipeline is developed to minimize redundant computations and enhance data reuse during rendering. The architecture is specifically designed to support neural implicit representations such as Neural Radiance Fields, enabling efficient sampling, interpolation, and neural network inference. Specialized processing units are integrated to accelerate these operations and achieve real-time performance. However, the coprocessor is mainly optimized for specific neural rendering tasks and may offer limited flexibility for broader graphics or computer vision applications. Hardware implementation also increases system complexity and cost. Furthermore, adapting the architecture to future neural rendering models may require additional modifications. Overall, the work highlights the effectiveness of specialized hardware acceleration in addressing computational and memory bottlenecks in neural rendering systems.[6]

Uni-Render: A Unified Accelerator for Real-Time Rendering Across Diverse Neural Renderers (2025), proposed by Chaojian Li, Sixu Li, Linrui Jiang, Jingqun Zhang, and Yingyan (Celine) Lin, introduces a flexible hardware accelerator designed to support multiple neural rendering methods in real time. The main objective of the work is to overcome the limitations of model-specific accelerators by providing a unified architecture for different neural rendering techniques. The proposed Uni-Render framework identifies common computational patterns across renderers such as Neural Radiance Fields (NeRF), point-based, and voxel-based methods. It integrates modular processing units, optimized memory hierarchies, and configurable dataflow mechanisms to accelerate ray sampling, feature processing, and neural inference. Hardware–software co-design techniques are also applied to improve performance and energy efficiency. However, supporting diverse rendering models increases architectural complexity. In addition, the unified architecture may not achieve maximum performance for a single renderer compared to specialized accelerators. Adapting the system to future neural rendering techniques may also require further architectural modifications. Overall, Uni-Render demonstrates that a unified accelerator approach can improve scalability and practicality for real-time neural rendering systems.[7]

Lumina: Real-Time Mobile Neural Rendering by Exploiting Computational Redundancy (2025), proposed by Yu Feng and colleagues, introduces an efficient neural rendering framework designed for mobile devices with strict power and resource constraints. The main goal of the study is to enable real-time neural rendering on mobile platforms while reducing computational overhead and maintaining visual quality. The proposed Lumina system identifies and removes redundant computations across rays, frames, and neural network layers using lightweight caching and reuse strategies. Optimized execution pipelines are designed to reuse intermediate results and minimize repeated neural inference. This approach significantly improves rendering speed and energy efficiency on mobile hardware. However, the effectiveness of the system depends on temporal and spatial coherence within scenes. Performance improvements may decrease in highly dynamic environments with frequent changes. In addition, the framework requires additional memory for caching intermediate results. Performance may

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

also vary across different mobile hardware platforms. Overall, Lumina demonstrates that exploiting computational redu]ndancy is an effective strategy for enabling real-time neural rendering on mobile and edge devices.[8]

PSA-NeRF: Personalized Spatial Attention Neural Rendering for Audio-Driven Talking Portraits Generation (2025), proposed by Huiyu Xu and colleagues, presents a neural rendering framework for generating realistic and personalized talking portrait videos from audio input. The main objective of the work is to improve lip synchronization, facial expression accuracy, and identity preservation in audio-driven portrait generation. The authors introduce PSA-NeRF, which extends Neural Radiance Fields by incorporating a personalized spatial attention mechanism that focuses on important facial regions such as the lips and jaw. Audio features are used to guide facial motion generation, while a subject-specific 3D representation ensures visual consistency across frames. This targeted attention improves audio–visual correspondence and enhances the realism of generated portraits. However, the approach requires person-specific training data, which limits scalability. Both training and rendering processes are computationally intensive. In addition, performance may decrease when audio inputs are noisy or highly expressive. Overall, PSA-NeRF demonstrates that integrating personalized spatial attention with NeRF can significantly improve audio-driven talking portrait generation for applications such as virtual avatars and digital humans.[9]

NeuHMR: Neural Rendering-Guided Human Motion Reconstruction (2025), proposed by Tiange Xiang and colleagues, presents a framework that utilizes neural rendering to improve the accuracy of 3D human motion and pose reconstruction from visual inputs. The main objective of the work is to reduce ambiguities in motion estimation by incorporating rendering-based supervision into the reconstruction process. The proposed NeuHMR framework integrates a neural rendering module with a motion reconstruction pipeline. Predicted 3D human body parameters are rendered into synthetic visual representations and compared with input observations to iteratively refine motion estimates. This joint optimization improves reconstruction accuracy and helps handle challenging conditions such as occlusions, complex poses, and viewpoint variations. However, the framework depends on accurate body models and high-quality training data. The integration of neural rendering and optimization also increases computational complexity. As a result, real-time deployment may require specialized hardware. Overall, NeuHMR demonstrates that neural rendering can effectively guide human motion reconstruction and improve robustness in motion analysis tasks.[10]

DNRSelect: Active Best View Selection for Deferred Neural Rendering (2025), proposed by Dongli Wu, Haochen Li, and Xiaobao Wei, introduces an approach to improve the efficiency and visual quality of deferred neural rendering systems. The main objective of the work is to reduce redundant rendering computations while maintaining high-quality novel view synthesis. The authors propose DNRSelect, an active view selection framework integrated into a deferred neural rendering pipeline. Multiple candidate viewpoints are evaluated using a learned scoring mechanism that estimates each view's contribution to the final output. Only the most informative views are selected, rendered, and fused, enabling more efficient use of computational resources. However, the effectiveness of the method depends on the accuracy of the view selection strategy. Incorrect or suboptimal view choices may lead to missing details or visual artifacts. The framework also introduces additional computational overhead for evaluating candidate views. Overall, DNRSelect demonstrates that intelligent view selection can significantly improve the efficiency of deferred neural rendering while preserving visual fidelity.[11]

MetaSapiens: Real-Time Neural Rendering with Efficiency-Aware Pruning and Accelerated Foveated Rendering, proposed by Weikai Lin, Yu Feng, and Yuhao Zhu, presents an efficient neural rendering framework for high-quality real-time human avatar rendering. The primary objective of the paper is to reduce computational cost while maintaining visual fidelity, particularly for immersive applications such as virtual reality and telepresence. The authors introduce MetaSapiens, which combines efficiency-aware model pruning with accelerated foveated rendering to optimize neural inference. Neural components with lower visual impact are pruned based on their contribution to perceptual quality, while foveated rendering allocates higher computational resources to regions near the viewer's gaze. This hybrid strategy significantly reduces rendering cost without noticeable degradation in perceived visual quality, enabling real-time performance. However, the framework depends on accurate gaze tracking for effective foveated rendering, and pruning strategies may require re-tuning for different avatar types or scene conditions. Performance may also degrade in scenarios involving rapid viewpoint or gaze changes. Overall, the paper demonstrates that integrating pruning and foveated rendering provides a practical and effective approach for real-time neural avatar rendering, advancing the deployment of neural rendering systems in immersive environments.[12]

Renderformer: Transformer-based Neural Rendering of Triangle Meshes with Global Illumination, proposed by Chong Zeng et al., introduces a neural rendering framework designed to render triangle mesh–based scenes with high visual quality while accurately modeling global illumination.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

The main objective of the paper is to overcome the limitations of traditional rasterization and point-based neural rendering by leveraging transformer architectures for flexible and realistic lighting representation. The authors propose Renderformer, which directly operates on triangle mesh primitives by encoding their geometric and material properties and employing a transformer to model long-range light interactions. Through attention-based relationships among mesh elements, the framework effectively captures complex lighting effects such as indirect illumination, soft shadows, and reflections. However, the transformer-based design incurs high computational and memory costs, particularly for dense or complex meshes, and requires large training datasets and powerful hardware. Achieving real-time performance for highly detailed scenes remains challenging without further optimization or specialized acceleration. Overall, the paper demonstrates that transformer-based neural rendering is a promising approach for mesh-based global illumination, marking an important step toward realistic and flexible neural graphics pipelines.[13]

NeuRSS: Enhancing AUV Localization and Bathymetric Mapping With Neural Rendering for Sidescan SLAM, proposed by Yiping Xie, Jun Zhang, Nils Bore, and John Folkesson, presents a neural rendering–based approach to improve underwater robot localization and mapping. The primary objective of the paper is to enhance Autonomous Underwater Vehicle (AUV) localization and bathymetric mapping by addressing challenges such as noisy sidescan sonar data and limited underwater visibility. The authors introduce NeuRSS, a framework that learns a neural representation of the relationship between 3D seafloor geometry and sidescan sonar observations. Neural rendering is used to synthesize realistic sonar images from estimated maps, which are then incorporated into a SLAM optimization process to refine both localization and mapping. This integration improves robustness to sonar artifacts and enhances map consistency. However, the approach requires substantial training data and computational resources, and performance may degrade in highly dynamic underwater environments or when sonar data quality is poor. Real-time deployment on resource-constrained AUV platforms also remains challenging. Overall, the paper demonstrates that neural rendering significantly enhances sidescan SLAM, contributing to more accurate and reliable underwater navigation and mapping.[14]

Real-time Neural Rendering of Dynamic Light Fields, proposed by Arno Coomans et al., presents a neural rendering framework designed to enable interactive rendering of dynamic light fields with high visual quality and low latency. The primary objective of the paper is to handle time-varying scenes and changing illumination while achieving real-time performance suitable for interactive applications. The authors introduce a neural approach that learns a compact spatio-temporal representation of dynamic light fields, allowing the model to predict novel views by jointly modeling spatial structure and temporal motion. Efficient network architectures and optimized data representations are employed to support fast inference and practical deployment in dynamic scenarios. However, the method requires substantial training data to accurately capture temporal variations, and rendering quality may degrade for highly complex or rapidly changing scenes. Additionally, memory and computational requirements can remain high for large-scale light field representations. Overall, the paper demonstrates that neural rendering is a viable solution for real-time dynamic light field rendering, contributing to the advancement of interactive neural graphics and real-time visualization systems.[15]

NeuRAD: Neural Rendering for Autonomous Driving, proposed by Adam Tonderski et al. from Zenseact, introduces a neural rendering framework specifically designed for autonomous driving applications. The primary objective of the paper is to generate realistic and sensor-consistent scene representations that support perception, simulation, and training of self-driving systems in complex real-world environments. NeuRAD adopts a learning-based neural scene representation approach tailored to automotive scenarios, modeling scene geometry, appearance, and view-dependent effects using multi-sensor driving data. The framework integrates neural rendering within a vehicle-centric coordinate system and is trained on large-scale driving datasets. It is evaluated through novel view synthesis and simulation tasks, demonstrating realistic road scene rendering and improved data efficiency for autonomous driving pipelines. However, the method requires large volumes of high-quality training data, and performance may degrade in rare or highly dynamic traffic scenarios. The computational cost of training is high, and real-time deployment may require specialized hardware acceleration. Overall, the paper demonstrates the strong potential of neural rendering for autonomous driving, marking an important step toward data-driven environment modeling for future self-driving technologies.[16]

Neural Rendering and Its Hardware Acceleration: A Review, authored by Xinkai Yan, Jieting Xu, Yuchi Huo, and Hujun Bao, provides a comprehensive overview of neural rendering techniques and the hardware acceleration requirements needed to support them efficiently. The primary objective of the paper is to analyze how deep learning integrates with traditional computer graphics and to identify the key challenges and future directions in neural rendering architectures. The authors conduct an extensive literature

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

review, categorizing neural rendering methods into forward rendering, inverse rendering, and post-processing applications. Various neural network models, including Multilayer Perceptrons (MLPs), Convolutional Neural Networks (CNNs), Generative Adversarial Networks (GANs), and Transformers, are examined in detail. The study further analyzes computational patterns, memory access behavior, and performance characteristics across different hardware platforms such as CPUs, GPUs, and domain-specific accelerators. However, neural rendering methods face challenges including high computational complexity, large memory demands, long training times, and limited generalization to unseen or dynamic scenes. Existing hardware platforms are also not fully optimized for neural rendering–specific operations like ray marching and hash-based encoding. Overall, the paper concludes that neural rendering is a promising direction for next-generation graphics applications, and emphasizes the need for algorithm–hardware co-design and dedicated accelerators to achieve real-time and energy-efficient performance.[17]

The Overview of Neural Rendering, authored by Zavalniuk Ye.K., Romanyuk O.N., Korobeinikova T.I., Titova N.V., and Romanyuk S.O., presents an analytical review of how neural networks enhance the efficiency and intelligence of image rendering processes. The primary objective of the paper is to examine the role of neural rendering techniques across various stages of three-dimensional scene visualization, including geometry generation, texture synthesis, lighting modeling, and image quality enhancement. The authors adopt a review-based methodology, analyzing different neural network architectures and their applications in rendering tasks such as triangular mesh generation, texture creation from photographs, texture style transfer, reflectance model representation, and reverse neural rendering for polygonal reconstruction from images. Despite their advantages, neural rendering approaches involve high computational complexity, require large training datasets, and demand significant hardware resources, making real-time implementation challenging. Accurate learning of complex lighting and reflectance models also remains difficult. Overall, the paper concludes that neural networks significantly improve the rendering pipeline by enabling intelligent processing at multiple visualization stages, contributing to more flexible, adaptive, and realistic image rendering systems.[18]

NEUREX: A Case for Neural Rendering Acceleration, authored by Junseo Lee, Kwanseok Choi, Jungi Lee, Seokwon Lee, Joonho Whangbo, and Jaewoong Sim, analyzes the computational challenges of neural rendering and presents NeuRex as a specialized acceleration framework to improve performance and energy efficiency. The primary objective of the paper is to highlight the limitations of general-purpose GPUs for neural rendering workloads and to demonstrate the necessity of hardware–software co-design. The authors conduct an architectural analysis of neural rendering tasks such as Neural Radiance Fields (NeRF), identifying key bottlenecks including irregular memory access, high computational demand, and inefficient data reuse. Based on this analysis, NeuRex is proposed as a domain-specific accelerator that optimizes ray sampling, neural network execution, and memory access patterns. Experimental evaluations show that NeuRex significantly outperforms GPU-based implementations in both performance and energy efficiency. However, the accelerator is primarily optimized for neural rendering tasks, limiting its applicability to other workloads, and the specialized hardware design increases development complexity and reduces flexibility for adapting to evolving models. Overall, the paper demonstrates that domain-specific acceleration is essential for real-time, high-quality neural rendering and strongly advocates hardware–software co-design in future graphics systems.[19]

NeRFahedron: A Primitive for Animatable Neural Rendering with Interactive Speed, proposed by Zackary P. T. Sin, Peter H. F. Ng, and Hong Va Leong, introduces a novel neural rendering primitive designed to support real-time animation and interactive performance in dynamic scenes. The primary objective of the paper is to overcome the limitations of existing neural rendering approaches that struggle with real-time speed and animatable scene representation. The authors propose NeRFahedron, a structured geometric primitive that combines neural scene representation with spatial decomposition, enabling efficient modeling of both appearance and geometry while supporting animation. Scenes are represented using multiple NeRFahedra integrated into a neural rendering pipeline, allowing smooth handling of dynamic content. Experimental evaluations demonstrate that the proposed framework achieves interactive rendering speeds while maintaining high visual quality. However, the method introduces additional complexity in scene decomposition and primitive management, and rendering quality depends on effective partitioning of scenes into NeRFahedra. Scalability to very large or highly complex scenes may also require careful tuning. Overall, the paper shows that NeRFahedron is an effective solution for animatable neural rendering at interactive speeds, contributing significantly toward practical real-time neural rendering systems for applications in interactive graphics, virtual reality, and animation.[20]

Advances in Neural Rendering(2022) authored by A. Tewari, J. Thies, B. Mildenhall, P. Srinivasan, E. Tretschk, W. Yifan,

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

C. Lassner, V. Sitzmann, R. Martin-Brualla, S. Lombardi, T. Simon, C. Theobalt, M. Nießner, J. T. Barron, G. Wetzstein, M. Zollhöfer, and V. Golyanik, presents a comprehensive overview of recent progress in neural rendering techniques that integrate computer graphics with deep learning to synthesize photorealistic images and videos. The primary objective of the paper is to systematically categorize existing neural rendering methods, analyze their fundamental principles, and identify open challenges and future research directions in neural scene representation and rendering. The authors adopt a structured literature review methodology, organizing approaches based on scene representation, rendering objectives, and learning strategies. Key techniques such as Neural Radiance Fields, deep voxel-based representations, point-based neural rendering, and hybrid graphics–learning pipelines are critically analyzed. The paper compares supervised and self-supervised training paradigms and evaluates methods using metrics including rendering quality, generalization capability, and computational efficiency. Despite significant advancements, neural rendering methods suffer from high computational costs, substantial memory requirements, and long training times. Many approaches struggle to achieve real-time performance and scalability for complex and dynamic scenes. Generalization to unseen viewpoints and varying lighting conditions also remains a challenge. Furthermore, dependence on large datasets limits practical deployment. The paper concludes that neural rendering has become a transformative paradigm in computer graphics, offering unprecedented realism and flexibility. However, further research is required to improve efficiency, robustness, and real-time capabilities. Overall, neural rendering is highlighted as a strong foundation for future applications in augmented reality, virtual reality, digital humans, and immersive media systems.[21]

This work presents a neural rendering framework for high-quality object reconstruction and novel view synthesis using large, unstructured online image collections. The primary objective is to overcome challenges such as sparse viewpoints, varying illumination, inconsistent backgrounds, and the absence of controlled capture conditions commonly found in internet images. The proposed approach jointly learns object geometry, appearance, and view-dependent effects through a learning-based neural rendering pipeline. Neural implicit representations are employed to model object shape and radiance, enabling realistic rendering across diverse viewpoints. A multi-stage training strategy is introduced to progressively refine reconstruction accuracy and visual quality, improving robustness across different object categories. The framework is evaluated using both qualitative and quantitative metrics, demonstrating strong performance under challenging real-world conditions. However, the method relies on the availability of sufficiently diverse image collections, and performance may degrade with extremely sparse, occluded, or background-dominated inputs. Training the model is computationally expensive, and extending the approach to dynamic or deformable objects remains challenging. Overall, the study demonstrates the feasibility of neural rendering from unconstrained online images and highlights its potential for scalable 3D object reconstruction and visualization in real-world scenarios.[22]

BokehMe: When Neural Rendering Meets Classical Rendering presents a hybrid approach for generating realistic and physically plausible bokeh effects by combining neural rendering techniques with classical optics-based rendering principles. The primary objective of the work is to overcome the limitations of purely data-driven neural methods and traditional depth-of-field rendering approaches by leveraging the strengths of both paradigms. The proposed framework integrates classical optical models to guide bokeh formation while employing neural networks to estimate scene depth, lens parameters, and refine fine image details. A learning-based pipeline predicts depth and optical characteristics, followed by a physically inspired rendering process that synthesizes high-quality bokeh effects. The system is trained and evaluated on diverse datasets, demonstrating robustness across different scenes and lighting conditions. However, the method relies heavily on accurate depth estimation, and performance may degrade in complex or ambiguous scenes. The hybrid design also increases computational cost and system complexity compared to purely classical approaches, and requires high-quality training data for strong generalization. Overall, the paper demonstrates that combining neural and classical rendering yields more realistic, controllable, and visually pleasing depth-of-field effects, highlighting the potential of hybrid rendering strategies for practical photography and image processing applications.[23]

This work proposes a neural rendering framework based on light field representation to enable realistic novel view synthesis from sparse input views. The primary objective is to overcome the limitations of traditional light field rendering, which requires dense viewpoint sampling, by using neural networks to efficiently model view-dependent appearance. The approach adopts a learning-based light field modeling strategy, where neural networks learn a continuous mapping from spatial and angular coordinates to radiance values. By integrating classical light field theory with deep learning, the framework enables smooth view interpolation and high-quality rendering from a limited number of input images. The model is trained end-to-end and evaluated on multiple datasets, demonstrating strong visual quality and consistency across viewpoints. However, accurate

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

reconstruction requires sufficient viewpoint coverage, and performance may degrade with extremely sparse, occluded, or complex scenes. Training the neural light field model is computationally expensive, and capturing high-frequency details remains challenging. Additionally, the current framework does not directly support dynamic scenes. Overall, the study shows that neural light field rendering is an effective and flexible solution for sparse-view novel view synthesis, with promising applications in virtual reality, augmented reality, and immersive visual media.[24]

This work introduces a generalizable neural rendering framework that enables novel view synthesis across diverse scenes using a patch-based representation. The primary objective is to overcome the limited generalization of scene-specific neural rendering methods by learning transferable local representations that can be applied to unseen environments. The proposed approach uses local image patches as fundamental rendering units instead of relying on a global scene representation. Neural networks are trained to predict view-dependent appearance from patch-level geometric and photometric features extracted from sparse input images. By focusing on local cues, the framework achieves strong generalization across different scenes and scene categories. Experimental evaluations on multiple datasets demonstrate high-quality novel view synthesis, robustness to varying scene content, and improved transferability compared to scene-specific methods. However, the patch-based strategy may struggle to maintain global scene consistency in highly complex environments. Rendering quality also depends on the availability of informative local patches, and performance may degrade with extremely sparse input views. Additionally, training requires significant computational resources and carefully curated datasets. Overall, the study shows that patch-based neural rendering is an effective and scalable approach for achieving scene generalization, making it promising for applications in augmented reality, robotics, and immersive visual systems.[25]

This work presents a neural rendering approach aimed at achieving high-quality, view-dependent novel view synthesis of real-world scenes from sparse multi-view image capture. The primary objective is to overcome the storage, scalability, and interpolation limitations of traditional light field and lumigraph techniques by introducing a compact neural representation suitable for immersive applications such as augmented reality, virtual reality, and free-viewpoint video. The proposed methodology integrates classical image-based rendering concepts with neural networks to model a continuous lumigraph function. A neural model is trained on sparsely sampled multi-view images to learn the relationship between camera viewpoints and pixel appearance, encoding spatial and angular information in a compact form rather than explicitly storing dense light fields. During rendering, the trained network synthesizes unseen viewpoints while preserving view-dependent effects such as reflections, specular highlights, and occlusions. However, the approach requires controlled multi-view capture setups and extensive preprocessing, and training the neural model is computationally expensive and time-consuming. The framework may also struggle with highly dynamic scenes, complex lighting variations, or viewpoints far outside the training distribution. Overall, the study demonstrates that compact neural representations can effectively replace dense lumigraph storage, marking a significant advancement in neural rendering for realistic and immersive visual experiences.[26]

FastNeRF addresses the major limitation of traditional Neural Radiance Fields, which produce high-quality renderings but suffer from extremely slow inference speeds that prevent real-time use. The primary objective of this work is to achieve high-fidelity neural rendering at real-time frame rates suitable for interactive applications such as virtual reality, augmented reality, and gaming. To accomplish this, the proposed FastNeRF framework redesigns the NeRF architecture by decoupling spatial and view-dependent components. Instead of evaluating a heavy multilayer perceptron at every sampled point along each ray, the method precomputes scene information into a structured feature grid that stores position-dependent geometry and appearance features. During rendering, efficient interpolation is used to retrieve spatial information, while a lightweight neural network models view-dependent effects such as specular highlights. This significantly reduces computational overhead during inference and enables rendering speeds of up to 200 frames per second while maintaining high visual quality. However, the approach requires an expensive preprocessing and training phase, increased memory consumption due to the structured grid, and is primarily designed for static scenes, limiting its applicability to dynamic environments. Overall, FastNeRF successfully bridges the gap between rendering speed and visual fidelity, representing an important advancement toward practical, real-time neural rendering systems.[27]

This work focuses on improving the representation of high-quality surface textures in volumetric neural rendering, which is a known limitation of existing neural rendering approaches that store appearance directly within volumetric density fields. The primary objective is to enhance texture fidelity, visual realism, and memory efficiency by decoupling texture representation from volumetric scene encoding. The proposed approach introduces a neural texture mapping framework that learns surface-aligned texture coordinates

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

and maps them to appearance features using a compact neural representation. Instead of embedding color information throughout a 3D volume, texture information is stored in a surface-aware manner and integrated into volumetric neural rendering pipelines during rendering. This design leads to sharper textures, better preservation of high-frequency details, and reduced memory consumption. The method is evaluated through extensive experiments and comparisons with existing volumetric neural rendering techniques, demonstrating improved rendering quality across different scenes and viewpoints. However, the approach depends on accurate surface geometry estimation, and performance may degrade in scenes with noisy or complex geometry. Training the neural texture representation is computationally expensive and requires careful tuning, and handling highly dynamic scenes remains challenging. Overall, this work shows that separating texture mapping from volumetric representation is an effective strategy for advancing neural rendering quality and efficiency, making it highly promising for applications in graphics, virtual reality, and 3D content creation.[28]

This work presents a neural rendering–based framework for automating the creation of realistic and controllable game characters, aiming to reduce the manual effort and expertise traditionally required in character modeling. The primary objective is to introduce a data-driven approach capable of generating high-quality character geometry, texture, and appearance for game development applications. The proposed methodology employs deep learning techniques to learn character attributes from input data such as images or captured character models. Neural networks are used to represent character geometry and appearance and to render consistent novel views. The framework integrates neural rendering with character parameterization, enabling automatic character generation and customization. Experimental evaluations demonstrate improved realism, visual consistency, and efficiency compared to conventional character creation pipelines. However, the approach relies heavily on the quality and diversity of training data, and performance may degrade for characters with uncommon poses, accessories, or complex clothing. Training the neural rendering models requires significant computational resources, and real-time deployment may require specialized hardware. Overall, the study shows that neural rendering is a promising solution for automated game character creation, offering scalable and efficient workflows for future interactive game content.[29]

This work introduces an equivariant neural rendering framework that explicitly incorporates geometric symmetries into neural scene representations to improve consistency and generalization. The main objective is to ensure that rendered outputs transform predictably under viewpoint changes such as rotations and translations. The proposed methodology is based on theory-driven network design using group-equivariant neural networks within the neural rendering pipeline. By enforcing equivariance constraints, the model learns structured representations aligned with 3D geometric transformations. This leads to improved robustness and data efficiency compared to conventional neural rendering approaches. The framework is evaluated on novel view synthesis tasks, where it demonstrates more stable and consistent rendering across different viewpoints. However, the approach introduces additional architectural complexity and computational overhead. Careful selection of symmetry groups is required to match specific rendering tasks. The method may also face challenges when applied to highly complex or non-rigid scenes. Overall, the study shows that embedding geometric priors directly into neural rendering models enhances stability and generalization, contributing to more reliable and principled neural rendering systems.[30]

3.OBJECTIVE

The main objective of this review paper is to analyze and understand the concept of neural rendering in computer graphics by bridging traditional rendering techniques with modern deep learning approaches. It systematically reviews around 30 research works published between 2020 and 2026, focusing on key methodologies such as Neural Radiance Fields (NeRF), transformer-based models, and light-field rendering. The paper aims to examine different neural rendering pipelines, scene representation methods, and hardware acceleration techniques including GPUs, NPUs, and specialized coprocessors. It also highlights major challenges such as high computational cost, latency, memory usage, and limited generalization of current models. Additionally, the review explores various application domains like autonomous driving, medical simulation, robotics, and virtual environments. Another important objective is to compare different research approaches and identify their strengths and limitations. Finally, the paper aims to provide insights into future research directions such as real-time rendering, model optimization, and hardware–software co-design to improve efficiency and scalability of neural rendering systems.

4. COMPARISON OF PAST PUBLISHED PAPERS

The comparison table is included to clearly summarize and compare different research papers in a structured format. It helps in understanding the key objectives, methodologies, and conclusions of each work quickly. It also highlights the differences and similarities between various approaches in neural rendering. Overall, it makes analysis easier and supports better understanding.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

Title of Paper	Year	Authors	Proposed Objective	Methodology	Conclusion
RenderFlow: Single-Step Neural Rendering via Flow Matching	2026	Zhang, S. et al.	Develop a fast, deterministic, physically consistent neural rendering framework addressing high latency of diffusion-based methods.	Single-step flow matching on a pretrained video Diffusion Transformer; G-buffer guidance; bridge matching in latent space.	Reformulating neural rendering as single-step flow matching efficiently bridges physically based and generative rendering.
Adaptive Neural Rendering for Dynamic ALS Team Training	2026	Alghamdy, B. S. et al.	Improve realism and responsiveness of medical training by dynamically adapting simulation environments.	Continuous volumetric neural rendering framework with adaptive learning mechanisms for real-time scene update.	Combining adaptive neural rendering with continuous volumetric representation enhances ALS simulation flexibility.
Real-Time Robot Self-Modeling via Direct Sensorimotor Decoding	2026	Asghar, M. Z.	Enable real-time robot self-modelling by eliminating neural rendering latency in robotic pipelines.	Direct sensorimotor decoding framework bypassing visual reconstruction; neural mapping of sensor inputs to model parameters.	Direct sensorimotor decoding overcomes rendering latency, enabling faster adaptation and more autonomous robotic systems.
Neural Scene Baking for Permutation Invariant Transparency Rendering	2026	Zhang, Z. & Simo-Serra, E.	Enable real-time rendering of transparent objects with global illumination, eliminating order-dependency artefacts.	Neural scene baking precomputes GI effects; permutation-invariant architecture ensures consistent transparent surface rendering.	Neural baking with permutation-invariant design provides a practical real-time solution for transparency with global illumination.
An Overview of Neural Rendering Accelerators	2025	Ryu, J. & Yoo, H.-J.	Bridge the gap between neural rendering algorithms and efficient hardware implementations.	Survey-based review of GPU, NPU, and dedicated accelerators; analysis of compression and co-design optimisations.	Specialised accelerator designs tailored to neural rendering workloads are necessary for real-time performance.

5. CONCLUSION

This review has examined 30 key publications in neural rendering spanning 2020 to 2026, revealing a field that has evolved rapidly from the foundational NeRF paradigm to a rich ecosystem of specialised techniques. The surveyed works collectively demonstrate that neural rendering has achieved photorealistic image synthesis and novel view generation with unprecedented flexibility, enabling applications ranging from interactive computer graphics to autonomous driving and medical simulation.

Several cross-cutting themes emerge from this review. First, the tension between rendering quality and computational efficiency remains the central challenge; works such as FastNeRF, Lumina, and MetaSapiens address this through architectural innovations including feature caching, redundancy elimination, and foveated rendering. Second, hardware acceleration is increasingly recognised as essential, with domain-specific accelerators (NeuRex, Uni-Render, neural rendering coprocessors) demonstrating order-of-magnitude improvements over general-purpose GPU baselines.

Third, generalisation beyond training distributions remains an open problem. Patch-based and equivariant approaches take steps toward scene-agnostic rendering, but performance in highly dynamic or unconstrained real-world settings still

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

lags behind controlled benchmarks. Fourth, domain-specific adaptations—for robotics, medical training, autonomous vehicles, and underwater mapping—demonstrate that neural rendering principles transfer broadly, but each domain introduces unique data constraints, latency requirements, and safety considerations.

In summary, neural rendering has matured from a research curiosity into a practical technology with demonstrated real-world impact. Continued advances in model compression, hardware–software co-design, and generalisation will be critical for the next generation of applications.

6. Future Scope

The trajectory of neural rendering research points toward several promising future directions that are likely to shape the field over the next decade:

6.1 Real-Time and Edge Deployment

While works such as FastNeRF and Lumina have made significant strides, achieving true real-time neural rendering on consumer and edge hardware (smartphones, AR glasses, embedded systems) remains an active goal. Future research should focus on aggressive model compression, neural architecture search for hardware-aware design, and quantisation-aware training tailored to low-power accelerators.

6.2 Dynamic and Deformable Scene Rendering

Most current methods assume static or quasi-static scenes. Extending neural rendering to fully dynamic environments—including deformable objects, fluid simulations, and crowd animations—will require new spatio-temporal representations, efficient temporal caching strategies, and physics-informed neural priors.

6.3 Hardware–Software Co-Design

The success of early neural rendering accelerators suggests that co-designing algorithms and hardware from the ground up will yield significant efficiency gains. Future work should explore programmable dataflow architectures, in-memory computing for neural inference, and standardised hardware APIs for neural rendering primitives.

6.4 Generalised and Foundation Models for Rendering

Large-scale pretrained generative models (e.g., diffusion transformers) are beginning to be adapted for rendering tasks, as seen in RenderFlow. Future research may leverage foundation models trained on massive 3D datasets to achieve universal, scene-agnostic rendering with minimal per-scene fine-tuning.

6.5 Integration with Physical Simulation

Bridging neural rendering with physics engines will unlock richer simulations for training autonomous systems, validating engineering designs, and creating interactive virtual worlds. Differentiable rendering pipelines that allow gradient flow through physical simulation are a particularly promising avenue.

6.6 Ethical and Privacy Considerations

As neural rendering enables hyper-realistic avatar generation and scene synthesis, ethical challenges around deepfakes, digital identity, and data privacy will intensify. Future research must address robust detection of synthetic media, consent-aware rendering systems, and standards for responsible deployment of neural rendering technologies.

REFERENCES

[1] Zhang, S., Liu, R., Schroers, C., & Zhang, Y. (2026). RenderFlow: Single-step neural rendering via flow matching. Ar Xiv preprint.
[2] Alghamdy, B. S., Alasmari, W. A., Alsiwed, D. T. S., Alalmaei, N. A., Alghamdi, E. A., Ghazal, H., & Almeshal, M. A. (2026). Adaptive neural rendering for dynamic advanced life support team training. Simulation in Healthcare.
[3] Asghar, M. Z. (2026). Real-time robot self-modeling via direct sensorimotor decoding: Overcoming the neural rendering latency bottleneck. arXiv preprint.
[4] Zhang, Z., & Simo-Serra, E. (2026). Neural scene baking for permutation invariant transparency rendering with real-time global illumination. ACM Transactions on Graphics.
[5] Ryu, J., & Yoo, H.-J. (2025). An overview of neural rendering accelerators: Challenges, trends, and future directions. IEEE Transactions on Circuits and Systems for Video Technology.
[6] Yuan, Z., Yuan, B., Rao, C., Zhu, Y., He, Y., Zhou, P., Yu, J., & Lou, X. (2025). A neural rendering coprocessor with optimized ray representation and marching. IEEE Journal of Solid-State Circuits.
[7] Li, C., Li, S., Jiang, L., Zhang, J., & Lin, Y. C. (2025). Uni-Render: A unified accelerator for real-time rendering across diverse neural renderers. Proceedings of ISCA.
[8] Feng, Y., Lin, W., Cheng, Y., Liu, Z., Leng, J., Guo, M., Chen, C., Sun, S., & Zhu, Y. (2025). Lumina: Real-time mobile neural rendering by exploiting computational redundancy. Proceedings of ASPLOS.
[9] Xu, H., Jin, S., Wang, Z., Ba, Z., & Wei, T. (2025). PSA-NeRF: Personalized spatial attention neural rendering for audio-driven talking portraits generation. Ar Xiv preprint.
[10] Xiang, T., Wang, K.-C., Heo, J., Adeli, E., Yeung, S., Delp, S., & Fei-Fei, L. (2025). NeuHMR: Neural rendering-guided human motion reconstruction. Proceedings of CVPR.
[11] Wu, D., Li, H., & Wei, X. (2025). DNRSelect: Active best view selection for deferred neural rendering. arXiv preprint.
[12] Lin, W., Feng, Y., & Zhu, Y. (2025). MetaSapiens: Real-time neural rendering with efficiency-aware pruning and accelerated foveated rendering. Proceedings of ISCA.
[13] Zeng, C., Dong, Y., Peers, P., Wu, H., & Tong, X. (2025). Renderformer: Transformer-based neural rendering of triangle meshes with global illumination. ACM SIGGRAPH.
[14] Xie, Y., Zhang, J., Bore, N., & Folkesson, J. (2025). NeuRSS: Enhancing AUV localization and bathymetric mapping with neural rendering for sidescan SLAM. IEEE Robotics and Automation Letters.
[15] Coomans, A., Dominici, E. A., Döring, C., Mueller, J. H., Hladky, J., & Steinberger, M. (2025). Real-time neural rendering of dynamic light fields. Proceedings of CVPR.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

[16] Tonderski, A., Lindström, C., Hess, G., Ljungbergh, W., Svensson, L., & Petersson, C. (2024). NeuRAD: Neural rendering for autonomous driving. Proceedings of CVPR.

[17] Yan, X., Xu, J., Huo, Y., & Bao, H. (2024). Neural rendering and its hardware acceleration: A review. ACM Computing Surveys.

[18] Zavalniuk, Y. K., Romanyuk, O. N., Korobeinikova, T. I., Titova, N. V., & Romanyuk, S. O. (2023). The overview of neural rendering. Journal of Applied Computer Science.

[19] Lee, J., Choi, K., Lee, J., Lee, S., Whangbo, J., & Sim, J. (2023). NEUREX: A case for neural rendering acceleration. Proceedings of MICRO.

[20] Sin, Z. P. T., Ng, P. H. F., & Leong, H. V. (2023). NeRFahedron: A primitive for animatable neural rendering with interactive speed. Proceedings of SIGGRAPH Asia.

[21] Tewari, A., Thies, J., Mildenhall, B., Srinivasan, P., Tretschk, E., Yifan, W., Lassner, C., Sitzmann, V., Martin-Brualla, R., Lombardi, S., Simon, T., Theobalt, C., Nießner, M., Barron, J. T., Wetzstein, G., Zollhöfer, M., & Golyanik, V. (2022). Advances in neural rendering. ACM Computing Surveys, 54(8), 1–38. https://doi.org/10.1145/3451256

[22] Kuang, Z., Olszewski, K., Chai, M., Huang, Z., Achlioptas, P., & Tulyakov, S. (2022). NeROIC: Neural rendering of objects from online image collections. ACM Transactions on Graphics, 41(4).

[23] Peng, J., Cao, Z., Luo, X., Lu, H., Xian, K., & Zhang, J. (2022). BokehMe: When neural rendering meets classical rendering. Proceedings of CVPR.

[24] Suhail, M., Esteves, C., Sigal, L., & Makadia, A. (2022). Light field neural rendering. Proceedings of CVPR.

[25] Suhail, M., Esteves, C., Sigal, L., & Makadia, A. (2022). Generalizable patch-based neural rendering. Proceedings of ECCV.

[26] Kellnhofer, P., Jebe, L. C., Jones, A., Spicer, R., Pulli, K., & Wetzstein, G. (2021). Neural lumigraph rendering. Proceedings of CVPR.

[27] Garbin, S. J., Kowalski, M., Johnson, M., Shotton, J., & Valentin, J. (2021). FastNeRF: High-fidelity neural rendering at 200FPS. Proceedings of ICCV.

[28] Xiang, F., Xu, Z., Hašan, M., Hold-Geoffroy, Y., Sunkavalli, K., & Su, H. (2021). NeuTex: Neural texture mapping for volumetric neural rendering. Proceedings of CVPR.

[29] Shi, T., Zou, Z., Shi, Z., & Yuan, Y. (2020). Neural rendering for game character auto-creation. Proceedings of AAAI.

[30] Dupont, E., Bautista, M. A., Colburn, A., Sankar, A., Susskind, J., & Shan, Q. (2020). Equivariant neural rendering. Proceedings of ICML.