

Distributed Radiance Field Training for 6G-Enabled Metaverse: Requirements and Challenges

Yendoubouam Lalle

Doctoral Candidate

AI Research Lab, Nokia Bell Labs

Massy, France

Fabio Pianese

Ph.D. Advisor

AI Research Lab, Nokia Bell Labs

Massy, France

Abstract—Immersive metaverse applications in future 6G networks will require efficient three-dimensional (3D) scene reconstruction. Techniques such as Neural Radiance Fields and 3D Gaussian Splatting offer high-quality representations but rely on centralized training with significant computational and communication demands. Distributed and federated approaches have been proposed to address these limitations by enabling decentralized processing across edge networks. This paper examines recent research on distributed radiance field training, highlighting open challenges such as communication overhead, client heterogeneity, and model synchronization. We outline possible directions for algorithmic improvement and propose architectural considerations aligned with 6G capabilities, including edge–cloud coordination and adaptive resource management. The aim is to clarify current limitations and contribute to a better understanding of the requirements for scalable and reliable scene understanding.

Index Terms—6G networks, radiance fields, federated learning, metaverse, distributed computing

I. INTRODUCTION

The metaverse envisions a persistent, immersive digital environment blending virtual and physical spaces, with applications spanning augmented reality, digital twins, and autonomous systems. Delivering such experiences at scale requires accurate and efficient 3D scene reconstruction, often performed across distributed sensors and edge devices. Techniques such as Neural Radiance Fields (NeRF) [1] and 3D Gaussian Splatting (3D-GS) [2] have achieved impressive results in photorealistic view synthesis. However, these methods rely on centralized data collection and training, which limits scalability and adaptability in dynamic or resource-constrained environments, such as vehicular networks or Ambient IoT contexts.

The emergence of 6G introduces new opportunities for distributed learning, with anticipated support for high bandwidth, low latency, and integrated edge intelligence. These capabilities are promising for offloading computation and enabling collaborative scene reconstruction and update. Nonetheless, applying NeRF or 3D-GS in such settings remains challenging. Centralized training struggles with the volume of data involved and lacks the flexibility required to adapt to changing network conditions or dynamic membership scenarios.

To address these constraints, recent research has explored distributed and federated training approaches. Techniques such as federated NeRF [3] and federated 3D-GS (Fed3DGS) [4]

propose training radiance fields across multiple devices while reducing communication overhead and preserving data privacy. While these methods show promise, they also introduce new challenges related to communication efficiency, handling heterogeneous or asynchronous clients, and maintaining model quality across distributed nodes.

This work seeks to clarify these open challenges and explore possible directions for improvement. In particular, we focus on three guiding questions:

- 1) How can radiance fields be partitioned and trained across heterogeneous edge nodes under bandwidth and compute constraints?
- 2) What 6G network capabilities (e.g., latency, throughput, slicing) and resource-sharing techniques are needed to support distributed learning and aggregation?
- 3) How can incremental model updates be managed to maintain scene fidelity under dynamic conditions?

In this extended abstract, we present a preliminary analysis of the algorithmic and architectural requirements for distributed radiance field rendering in 6G-enabled environments. We discuss current limitations in the literature, identify potential areas for algorithmic improvement such as partitioning strategies, model compression, and asynchronous aggregation, and outline architectural considerations relevant to future device–edge–cloud systems. The rest of this paper is organized as follows: Section II reviews relevant background on radiance field methods, distributed training approaches, and 6G communication frameworks; Section III presents a set of research challenges and outlines possible directions for algorithmic and architectural advancement; Section IV illustrates these ideas through a vehicular metaverse use case; Section V concludes the paper and outlines future research directions.

II. BACKGROUND AND RELATED WORK

A. Neural Radiance Fields and 3D Gaussian Splatting

Radiance fields are neural scene representations that enable photorealistic novel view synthesis from sparse input images. Two major families of approaches have emerged: Neural Radiance Fields (NeRF) and 3D Gaussian Splatting (3D-GS), which differ significantly in how scenes are modeled, trained, and rendered.

NeRF is an implicit representation, where a multilayer perceptron (MLP) encodes the scene as a continuous volumetric function of 3D coordinates and view directions [1]. Rendering is performed by querying the MLP along rays and accumulating color and density through volume rendering. Training a NeRF involves optimizing network weights to fit the appearance of a scene based on input images and camera poses, typically requiring many iterations and substantial compute resources. Although NeRF models are relatively lightweight in memory, they are computationally intensive to train and evaluate, and the monolithic MLP structure complicates distributed training.

In contrast, 3D-GS [2] is an explicit representation based on a collection of 3D Gaussian primitives, each with position, size, orientation, opacity, and color attributes. These Gaussians are projected and rasterized directly to render images, allowing fast and parallel rendering on GPUs [5]. Training involves optimizing the properties of these Gaussians to match the scene appearance. While 3D-GS models can be larger than NeRFs in memory, they support real-time rendering and are more amenable to editing or merging, making them attractive for distributed and interactive applications. From a distributed deployment perspective, NeRF's MLP structure complicates model partitioning and synchronization, whereas 3D-GS's spatial structure may offer better opportunities for scene decomposition. However, 3D-GS models can be storage-intensive and still require large-scale optimization across many parameters [5].

B. Federated and Distributed Learning for Radiance Fields

Recent work has begun to decentralize radiance field training to meet the demands of edge and 6G environments. Table I summarizes representative methods and their primary limitations, including Federated NeRF (FedNeRF), Federated 3D Gaussian Splatting (Fed3DGS), Decentralized NeRFs (DecentNeRFs), Distributed NeRF (Di-NeRF), and Distributed Optimization for Gaussian Splatting (DOGS).

Federated NeRF variants represent the first systematic approach to distributed radiance field training, with initial work by Holden et al. [3] and subsequent enhancements by Suzuki [6] introducing privacy-preserving aggregation mechanisms and client selection strategies to handle heterogeneous data distributions. Fed3DGS [4] extends federated learning to explicit Gaussian representations, using a distillation-based model update scheme to reconcile non-IID local data. While this approach benefits from the spatial structure of 3D Gaussians, the communication burden remains high due to the large number of parameters inherent in Gaussian splatting models.

DecentNeRFs [7] further advance the federated paradigm by separating private and global scene components, greatly reducing server load but introducing challenges in consistency when merging heterogeneous updates across diverse user environments. In contrast, Di-NeRF [8] adopts a peer-to-peer mesh topology where each agent refines both its local NeRF and relative poses with neighbors, supporting collaborative scene mapping without centralized coordination. While robust

to pose estimation errors, this approach demands reliable inter-agent communication links and sophisticated synchronization protocols that may be challenging in dynamic network conditions.

Recent advances in distributed Gaussian optimization, such as DOGS [9], explore spatial decomposition strategies that partition scene representations across multiple nodes. These methods show promise for scalable training but face limitations in handling increasing numbers of participants and maintaining pose alignment across distributed contributors.

Each method addresses different aspects of decentralization: FedNeRF focuses on privacy preservation, Fed3DGS tackles communication efficiency, DecentNeRFs handles user diversity, Di-NeRF addresses pose alignment, and DOGS explores spatial decomposition. However, significant challenges remain in achieving scalable, asynchronous radiance field training over heterogeneous edge networks. In the next section, we build on these insights to propose algorithmic improvements and architectural guidelines tailored for 6G-enabled metaverse systems.

III. DISTRIBUTED ARCHITECTURE AND ALGORITHMIC ENHANCEMENTS

A. Algorithmic Improvements for Distributed Training

Training radiance fields in distributed environments introduces several algorithmic challenges, including data and model partitioning, communication efficiency, and robustness to network heterogeneity. Addressing these requires rethinking how training is orchestrated across diverse and resource-constrained edge nodes. A key strategy is scene-aware model partitioning. In NeRF-based systems, the MLP can be segmented either layer-wise or by embedding dimensions, allowing each edge device to train a lightweight sub-model locally. In contrast, 3D-GS models can be spatially divided into clusters of Gaussians, which are then assigned to specific edge nodes for local refinement. This partitioning minimizes redundant computation and communication by limiting updates to relevant spatial or functional subsets [3], [4]. Another essential improvement is communication-efficient updates. Instead of transmitting full model parameters, clients may share only the most informative updates or compressed surrogates, such as distilled render outputs or low-dimensional embeddings [4]. Methods such as federated distillation enable exchanging these compressed representations rather than large model weights, significantly reducing uplink load while maintaining convergence. To handle the lack of coordination and device heterogeneity, aggregation schemes have been developed that integrate client updates upon arrival using staleness-aware weighting, thereby mitigating the impact of straggling nodes and unstable network links. In peer-to-peer setups such as Di-NeRF [8], collaborative training proceeds through mesh-based communication, where nodes refine their local models and relative poses based on local data and neighbor exchanges.

TABLE I
OVERVIEW OF FEDERATED/DISTRIBUTED RADIANCE FIELD METHODS

| Method | Key Contribution | Limitations |
|------------------|---|---|
| FedNeRF [3], [6] | Federated NeRF training with privacy-preserving aggregation and client selection strategies | Communication overhead; limited handling of highly heterogeneous scenes |
| Fed3DGS [4] | Applies federated learning to 3D Gaussian Splatting with distillation-based updates | Large parameter sizes; communication burden; handles only simple appearance shifts |
| DecentNeRFs [7] | Decomposes personal vs. global NeRFs; aggregates only shared content | Inconsistent integration across diverse users; coordination overhead |
| Di-NeRF [8] | Jointly optimizes NeRF parameters and relative poses in distributed mesh topology | Requires stable mesh connectivity; complex synchronization protocols |
| DOGS [9] | Distributed optimization for Gaussian splatting with spatial decomposition strategies | Limited scalability with increasing number of participants; pose alignment challenges |

B. Architectural Support in 6G Networks

The emergence of 6G opens new possibilities for deploying radiance field training over distributed infrastructures. A natural fit is a device–edge–cloud hierarchy, where lightweight local updates are performed on end devices (e.g., AR/VR headsets, connected vehicles), mid-level aggregation occurs at Multi-Access Edge Computing (MEC) nodes, and global consistency is maintained by cloud servers. This structure may alleviate communication bottlenecks and improve responsiveness for local scenes while enabling coarser-grained consistency management at higher levels [5]. Within this hierarchy, network slicing and AI-driven orchestration play a central role. Network slices tailored for metaverse workloads can provide guaranteed latency and bandwidth for periodic synchronization. AI-based orchestration agents deployed at MEC nodes can dynamically allocate compute resources, prioritize urgent updates, and optimize training rounds based on network load, energy profiles, and task criticality [10]. This tight integration between computation and communication layers is essential for real-time performance. Moreover, hierarchical aggregation enables context-aware optimization of resource allocation decisions. MEC nodes can maintain region-specific scene representations, updating only local segments of the global model. Such selective refinement, especially when combined with importance-driven update filtering, further improves scalability and reduces data movement—both crucial in multi-user, bandwidth-constrained scenarios.

C. 6G Infrastructure Requirements

For distributed radiance field rendering to be practically deployable, several infrastructural capabilities must be leveraged by future 6G networks. First, high-throughput wireless connectivity is essential. Even with model compression, exchanging updates at scale requires reliable uplink and downlink throughput on the order of hundreds of megabits per second per user [11]. This is particularly true for dense scenes with high-resolution geometry and appearance models. Second, ultra-reliable low-latency communication (URLLC) is needed to

minimize synchronization delays in training rounds. Realistic end-to-end latencies in the range of tens to hundreds of milliseconds for completing training updates are desirable to support responsive feedback loops in dynamic or collaborative scenarios [11]. Third, the availability of edge AI accelerators—including GPUs or neural processing units (NPUs) in vehicles, headsets, and MEC nodes—is a prerequisite for performing low-latency inference and training locally. This minimizes dependency on centralized infrastructure and allows fast model adaptation in situ. As 6G envisions the integration of AI at the edge, devices are expected to incorporate specialized hardware (such as NPUs) to support real-time 3D perception.

Finally, the ecosystem must support standardized orchestration APIs and protocols for model aggregation, versioning, client selection, and performance monitoring. These interfaces will be critical for ensuring interoperability across vendors and enabling secure federated learning at scale in heterogeneous 6G environments. For example, emerging 6G frameworks already emphasize management and orchestration specifications to handle slicing and service-level agreements, which could be extended to federated learning and distributed radiance-field operations.

IV. A VEHICULAR 6G METAVERSE USE CASE

A. Scenario and Methodology

In a representative vehicular metaverse scenario, connected vehicles and roadside infrastructure collaboratively maintain a photorealistic, continuously updated 3D digital twin of an urban area for both machine perception and human-facing AR services. Each vehicle captures multi-view imagery and sensor data, refining a local radiance-field model of its surroundings. Compact updates are periodically shared with nearby Multi-Access Edge Computing (MEC) servers, which aggregate regional models, manage cacheable radiance-field representations, and coordinate with cloud services for cross-region consistency. This device-edge-cloud hierarchy reflects 6G edge-intelligence design principles and aligns with recent proposals for distributed radiance-field modeling in autonomous driving [12], [13].

Several adaptations make radiance fields suitable for this distributed setting. *Spatial partitioning* can assign each node parameters for its region of interest, whether represented as Gaussian clusters or local neural embeddings. *Model compression and knowledge distillation* techniques can reduce bandwidth by exchanging compact representations such as distilled outputs or compressed parameter subsets [13]. To address mobility and dynamic environments, *incremental updates* can transmit only the most significant scene changes, reducing communication overhead in vehicular networks.

Before deployment, such pipelines could be prototyped using simulation platforms like CARLA [12], with evaluation based on standard metrics such as PSNR, SSIM, and task-specific performance measures.

B. Key Challenges

Several interrelated challenges remain for making such systems practical. *Communication efficiency* must be balanced against reconstruction fidelity, as aggressive compression can save bandwidth but degrade visual or perception quality [13]. *Heterogeneous client capabilities* and non-IID data distributions complicate model aggregation, requiring adaptive strategies for client participation and update weighting. *Pose estimation errors* and dynamic objects introduce alignment challenges in collaborative scene reconstruction [12]. *Intermittent connectivity* in urban V2X networks demands robust asynchronous aggregation and MEC-based coordination mechanisms. *Privacy and security* concerns arise when sharing scene representations, as distributed learning schemes face potential vulnerabilities to adversarial attacks. Additionally, *real-time constraints* in safety-critical driving scenarios require low-latency inference while maintaining high-quality scene reconstruction. Addressing these challenges will require integrated evaluation frameworks that jointly assess rendering fidelity, communication cost, latency, and system robustness under realistic 6G network conditions.

V. CONCLUSION AND FUTURE DIRECTIONS

This paper examined the emerging challenge of distributed radiance field training for metaverse applications in 6G environments. We reviewed recent research, identified core limitations such as high communication overhead, heterogeneous client capabilities, and the difficulty of maintaining model consistency, and outlined potential algorithmic and architectural directions to address them. The proposed concepts form a preliminary approach for developing scalable and efficient systems that can support immersive applications in future 6G networks. Key open questions include quantifying the trade-off between compression and rendering quality, characterizing convergence under heterogeneous and mobile clients, and ensuring secure and robust federated updates for safety-critical use cases. Furthermore, rapid advances in radiance field synthesis techniques may require technological adjustments to effectively incorporate emerging methods.

As next steps, we will develop prototype pipelines and controlled experiments that couple radiance field training

with realistic network and mobility conditions. We plan to define composite evaluation metrics that jointly capture visual fidelity, perception utility, bandwidth cost, and convergence behavior, and to iterate on partitioning, compression, and aggregation policies based on empirically observed trade-offs. In particular, we will employ machine learning techniques to identify viable technical trade-offs and to guide the design of evidence-driven algorithms and system architectures that are practical under real-world 6G deployment scenarios.

ACKNOWLEDGEMENTS

The work described in this document is funded by the European Commission's MSCA EXACT-6G doctoral network (grant agreement n. 101120297).

REFERENCES

- [1] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis," Aug. 2020, arXiv:2003.08934 [cs]. [Online]. Available: <http://arxiv.org/abs/2003.08934>
- [2] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis, "3D Gaussian Splatting for Real-Time Radiance Field Rendering," Aug. 2023, arXiv:2308.04079 [cs]. [Online]. Available: <http://arxiv.org/abs/2308.04079>
- [3] L. Holden, F. Dayoub, D. Harvey, and T.-J. Chin, "Federated Neural Radiance Fields," May 2023, arXiv:2305.01163 [cs]. [Online]. Available: <http://arxiv.org/abs/2305.01163>
- [4] T. Suzuki, "Fed3DGS: Scalable 3D Gaussian Splatting with Federated Learning," Mar. 2024, arXiv:2403.11460 [cs]. [Online]. Available: <http://arxiv.org/abs/2403.11460>
- [5] G. Wu, Z. Lyu, J. Zhang, and J. Xu, "Embracing Radiance Field Rendering in 6G: Over-the-Air Training and Inference with 3D Contents," Jun. 2024, arXiv:2405.12155 [cs]. [Online]. Available: <http://arxiv.org/abs/2405.12155>
- [6] T. Suzuki, "Federated learning for large-scale scene modeling with neural radiance fields," *arXiv preprint arXiv:2309.06030*, 2023.
- [7] Z. Tasneem, A. Dave, A. Singh, K. Tiwary, P. Vepakomma, A. Veeraraghavan, and R. Raskar, "DecentNeRFs: Decentralized Neural Radiance Fields from Crowdsourced Images," Mar. 2024, arXiv:2403.13199 [cs]. [Online]. Available: <http://arxiv.org/abs/2403.13199>
- [8] M. Asadi, K. Zareinia, and S. Saeedi, "Di-NeRF: Distributed NeRF for Collaborative Learning with Relative Pose Refinement," Sep. 2024, arXiv:2402.01485 [cs]. [Online]. Available: <http://arxiv.org/abs/2402.01485>
- [9] Y. Chen and G. H. Lee, "Dogs: Distributed-oriented gaussian splatting for large-scale 3d reconstruction via gaussian consensus," *Advances in Neural Information Processing Systems*, vol. 37, pp. 34487–34512, 2024.
- [10] Y. Zhang and Z. Shao, "Federated Neural Radiance Field for Distributed Intelligence," Jun. 2024, arXiv:2406.10474 [cs]. [Online]. Available: <http://arxiv.org/abs/2406.10474>
- [11] H. Pennanen, T. Hänninen, O. Tervo, A. Tölli, and M. Latva-aho, "6G: The Intelligent Network of Everything," *IEEE Access*, vol. 13, pp. 1319–1421, 2025, arXiv:2407.09398 [cs]. [Online]. Available: <http://arxiv.org/abs/2407.09398>
- [12] M. Dopirak, J. Gerec, and J. Gazda, "Reconstruction of Photorealistic 3D Urban Scenes Using Radiance Fields as Digital Twins for Autonomous Driving," *Acta Electrotechnica et Informatica*, vol. 24, no. 4, pp. 27–34, Dec. 2024. [Online]. Available: <https://www.sciendo.com/article/10.2478/aei-2024-0015>
- [13] E. Šlapak, M. Dopirak, M. A. A. Faruque, J. Gazda, and M. Levorato, "Distributed Radiance Fields for Edge Video Compression and Metaverse Integration in Autonomous Driving," in *2024 IEEE International Conference on Smart Computing (SMARTCOMP)*, Jun. 2024, pp. 71–76, arXiv:2402.14642 [cs]. [Online]. Available: <http://arxiv.org/abs/2402.14642>