Journal of Technological Information, Management & Engineering Sciences (2024),5(2), 1-15

Optimizing Real-Time Rendering Performance Through Adaptive Level-of-Detail Techniques

[1] **Dr. Saad Hassan**
[2] **Hira Mahmood**

How to cite this article:

Hassan, S., & Mahmood, H. (2024). Optimizing Real-Time Rendering Performance Through Adaptive Level-of-Detail Techniques. *Journal of Technological Information, Management & Engineering Sciences, 5(2), 1–15.*

Received	Accepted	Published
10 July 2024	15 September 2024	10 November 2024

Abstract

When rendering large or geometrically complex scenes in real time, there are issues of visual fidelity, interactive frame rates, memory consumption and energy consumption. Adaptive level-of-detail (LOD) techniques tackle this problem by choosing, combining or creating simpler representations of the scene based on view-dependent and resource-dependent conditions. This paper takes a critical synthesis of adaptive LOD methods for real-time rendering that range from discrete and continuous polygonal simplification, hierarchical and view-dependent LODs, virtualized geometry, to emerging methods such as neural and Gaussian-splatting representations. The review reveals that there is a lack of research that offers a single, perceptually-driven policy that can adapt to distance, movement, screen-space error, occlusion, hardware capabilities, and dynamic workload, and that does so simultaneously. This paper introduces a conceptual design of such an adaptive LOD controller and presents an evaluation method of the controller by the frame time, memory bandwidth, geometric error, temporal stability and perceptual quality. Instead of providing new empirical data, the paper summarizes published performance data, and outlines design directions for future systems. The analysis reveals that scalable real-time rendering has three key elements: hierarchical LOD structure, screen-space error metrics, and continuous transitions, while neural representations bring novel opportunities and costs for adaptive detail management.

Keywords: *adaptive level of detail, real-time rendering, mesh simplification, hierarchical level of detail, neural rendering, 3D Gaussian splatting, performance optimization*

Introduction

[1] Department of Computer Science, Air University, Islamabad, Pakistan.
[2] Department of Computer Science, Bahria University, Islamabad, Pakistan.

Copyright (c) 2024 Author(s)
This work is licensed under a Creative Commons Attribution 4.0 International License.

Real-time graphics applications such as video games, virtual reality, digital twins, simulation and interactive visualization demand a rendering system to deliver believable graphics at a specific time limit. With the typical goal of 60 frames per second, it takes about 16.67 milliseconds to draw the entire frame, which includes drawing simulation, traversing scene, processing the geometry, rasterization/ray tracing, shading and the presentation. The number of polygons, textures, lights and volumetric samples that need to be processed at full resolution becomes too expensive as the scene complexity grows. The level of detail techniques are one way to reduce this problem, since more expensive detail renderings of an object can be replaced by less expensive approximations when that information is not likely to be significant to the final image.

It's not just a count of polygons that are being reduced by LOD. Classical formulations include viewing distance, screen space projection size, motion, eccentricity and task demands (Luebke, 2001; Luebke et al., 2003). Today's engines build on these concepts with virtualized geometry, hierarchical scene representations, streaming and perceptually driven simplification. Recent approaches also use LOD in the case of neural radiance fields, implicit signed distance functions, and 3D Gaussian splatting, where the unit of detail is not a triangle, but rather a neural feature cell, a volumetric block, or a collection of Gaussians respectively (Kerbl et al., 2024; Takikawa et al., 2021).

The main challenge of this paper is to tackle the question of how an adaptive LOD can provide the best real-time rendering performance without compromising perception of the quality. The paper makes three contributions. First, it examines the theoretical underpinning and the more important families of adaptive LOD techniques. Secondly, it introduces the concept of limitations in the existing approaches, particularly in relation to a lack of common adaptation policies, which take into account geometry, perception, and hardware constraints. Thirdly, it presents a future adaptive LOD system conceptual framework and evaluation methodology. Is a review based paper, an academic synthesis, does not introduce or report new experimental measurements.

Literature Review

[1] Department of Computer Science, Air University, Islamabad, Pakistan.
[2] Department of Computer Science, Bahria University, Islamabad, Pakistan.

Copyright (c) 2024 Author(s)
This work is licensed under a Creative Commons Attribution 4.0 International License.

Foundations of Level-of-Detail Rendering

LOD is based on the concept that geometric details are more important visually if they are large and in good lighting conditions. The development of the hierarchical geometric model was introduced by Clark (1976), where objects can be described at different levels of detail, and the rendering system chooses the level of detail to render depending on the location of the viewer. This concept proved to be the basis for interactive graphics, where the more important parts of the model could be depicted in detail, while other parts were shown as projections.

Later studies have codified the process of mesh simplification as a technique for building LOD representations. Hoppe (1996) introduced progressive meshes, a series of approximations that are produced by successive edge-collapse operations. This can be useful for streaming and view-dependent operations for refinement and coarsening. Garland and Heckbert (1997) developed quadric error metrics, which measure the geometric penalty for contracting vertex pairs, yielding good quality, simplified surfaces with reasonable computational and storage cost. The method continues to be a very influential process since it offers a clear and measurable error criterion for determining when the process of simplification is acceptable.

Luebke (2001) conducted a survey of polygonal simplification algorithms and classified them into three categories, static, view dependent, and perceptually driven. Static LOD: Pre-computes multiple versions of an object and chooses one at runtime. View-dependent LOD: different parts of a single object can have different resolutions, depending on the camera. Perceptual LOD takes advantage of human visual models such as eccentricity, velocity, and contrast sensitivity to distribute detail where it is likely to impact perception. These principles were summarized by Luebke et al. (2003) in a treatment of LOD for 3D graphics that focused on simplification, hierarchy, error metrics and perceptual modulation.

Hierarchical LOD and View-Dependent LOD

Hierarchical LOD (HLOD) is an extension of object-level substitution to the scene and scene partitions. A hierarchy arranges a set of spatially adjacent geometry

[1] Department of Computer Science, Air University, Islamabad, Pakistan.
[2] Department of Computer Science, Bahria University, Islamabad, Pakistan.

Copyright (c) 2024 Author(s)
This work is licensed under a Creative Commons Attribution 4.0 International License.

into nodes, where each node can have increasingly coarse representations. The system then traverses the hierarchy from coarser to finer and chooses nodes where the error projected over the coarse nodes is still within a certain bound. The strategy can be helpful in many cases, such as large scanned models, terrain, city scenes, and industrial assemblies, where only a limited amount of geometry needs to be traversed, and the resulting geometry has limited impact on the image.

In more recent HLOD research, two practical artifacts have been tackled: popping - an abrupt change in the visible structure of the scene when the renderer changes to a higher or lower detail level, and cracking - discontinuities between neighbouring simplified patches. A solution to both problems was proposed in 2023 through a new visualization scheme based on partitioning the models on a cubic grid, building an octree-like hierarchical structure and performing vertex interpolation between each level of the hierarchy in the rendering process (CVMP, 2023). This shows that adaptive LOD should not only consider geometric reduction, but also temporal and spatial continuity.

View-dependent simplification is particularly important for large environments where various parts of the environment take up different portions of the screen. A building in the distance can be shown as a low-resolution shell model, and a building in the foreground can have a fine geometry. But view-dependent techniques involve extra bookkeeping and seams can occur if neighboring regions use different resolutions. Therefore, the well-developed systems are based on spatial hierarchies, error bounds, and transition mechanisms, in order to sustain visual coherence.

Virtualized Geometry and Modern Game Engines.

Virtualised geometry is a modern industrial implementation of adaptive LOD. Instead of small number of separate meshes per object, the system considers geometry as an ordered flow of clusters that can be retrieved and rasterized based on their visibility and screen-space cost. The authors of Karis et al. (2021) introduced the virtualized geometry system called Nanite by Unreal Engine that they claim to be able to render extremely dense meshes at interactive frame rates,

[1] Department of Computer Science, Air University, Islamabad, Pakistan.
[2] Department of Computer Science, Bahria University, Islamabad, Pakistan.

Copyright (c) 2024 Author(s)
This work is licensed under a Creative Commons Attribution 4.0 International License.

by means of fine-grained clustered LOD and software rasterization. The method goes beyond the artist-managed asset pipeline and moves to an automated runtime LOD system, depending on the screen-space error.

Virtualized geometry is important because it eliminates the need to write LOD chains by hand, and allows for the creation of film-quality assets for interactive scenes. However, it poses some problems with memory management, streaming latency, cache coherence, and with traversal of hierarchical structures. These costs are significant when the platform has limited bandwidth and/or power, and the geometric detail that can be achieved is limited.

The paper proposes adaptive LOD in Neural and Gaussian Representations.

The idea of LOD has been extended beyond polygonal meshes with neural rendering. Takikawa et al. (2021) introduced an octree representation of neural signed-distance functions, Neural Geometric Level of Detail (NGLOD), which can be rendered in multiple discrete LODs and interpolated between LODs continuously. The method allows to render high fidelity implicit shapes at real time without traversing the whole octree, while significantly outperforming fixed-size neural representations in terms of efficiency.

LOD is also used in radiance-field methods for handling large scale scenes. Gu et al. (2023) divided a scene into sub-networks to be trained and rendered with meshes of different levels of detail (LOD) style. The system has neural rendering capabilities, along with Unreal Engine rasterization, and reports real-time 4K rendering at up to 43 frames per second for large-scale scenes. Song et al. (2024) introduced a new approach called City-on-Web, a hybrid method that uses volume rendering, LOD and dynamically loading and unloading resources, achieving around 32 fps for large-scale neural rendering in the web context on an RTX 3060 GPU.

Another key aspect of adaptive LOD is 3D Gaussian splatting. Kerbl et al. (2023) have presented 3D Gaussian splatting for real-time rendering of the radiance field with high quality novel-view synthesis rate of up to interactive. This work was extended by Kerbl et al. (2024) in a hierarchy of 3D Gaussians for very large datasets. Their approach is to train chunks of the scene separately, build a

[1] Department of Computer Science, Air University, Islamabad, Pakistan.
[2] Department of Computer Science, Bahria University, Islamabad, Pakistan.
Copyright (c) 2024 Author(s)
This work is licensed under a Creative Commons Attribution 4.0 International License.

hierarchy of them and choose LODs for contents located at a great distance while ensuring that a smooth transition is possible between two LODs. This is a reflection of a general trend – adaptive detail is being increasingly realised through hierarchical resource selection, not just through the usual polygons reduction.

The work has been expanded to composed scenes with many Gaussians in more recent work using cluster-based LOD. Developed offline, the hierarchical clusters in a SIGGRAPH 2025 system are used to select their corresponding clusters online based on a footprint tolerance to reduce the number of Gaussians that have to be processed while maintaining visual quality (Chen et al., 2025). Likewise, Haydel and Yuksel (2023) proposed locally adaptive LOD for hardware-accelerated ray tracing, with a view to minimize energy consumption and memory traffic by varying LOD levels per ray. These techniques illustrate that LOD is slowly emerging as an interweaved resource management mechanism for geometry, neural fields and ray-tracing workloads.

Research Gap

Even though there has been significant progress, there are a number of gaps in the literature. First, many systems rely on distance or projected size as the primary LOD criterion and neglect to consider task relevance, content saliency, foveation or motion. Theoretically supported perceptually guided LOD has a long history, but is not always part of modern virtualized-geometry and neural rendering pipelines. Second, modern approaches tend to optimize a specific representation type (triangles, implicit fields, radiance fields, or Gaussians), while in practice, more and more diverse assets are increasingly being combined. There are no studies on a unified controller that distributes detail among these representations. Third, adaptive LOD decisions are often considered as local geometry-selection issues instead of global resource-allocation issues. The number of polygons is not the only factor in frame time; the number of draw calls, the cache behavior, the memory bandwidth, the complexity of the shaders, the cost of the ray-traversal and the streaming latency are also a factor. Very few published frameworks

[1] Department of Computer Science, Air University, Islamabad, Pakistan.
[2] Department of Computer Science, Bahria University, Islamabad, Pakistan.
Copyright (c) 2024 Author(s)
This work is licensed under a Creative Commons Attribution 4.0 International License.

explicitly use LOD selection as a constrained optimization problem with time and memory budget per frame. Finally, evaluations are also quite variable, with some studies reporting frames per second, others reporting perceptual similarity, geometric error, or memory use, and comprehensive evaluations that simultaneously measure performance, temporal stability, and perceived quality are rarely found.

Hence, the research gap that is being filled by this paper is the lack of a unified perceptually-aware, resource constrained adaptive LOD framework that works for both traditional and neural scene representations. In response, the paper summarizes current efforts and outlines a conceptual framework and evaluation methodology for such a framework.

Objectives

The intent of this paper is to investigate the use of adaptive LOD approaches to render the scenes with sufficient visual fidelity and at the same time reduce the performance overheads. The objectives are as follows:

Understand the theoretical background and the main classes of adaptive LOD techniques for real-time rendering.

To understand the influence of Hierarchical structures, Screen space error metrics and Continuous transitions on scalable rendering.

Examining extending LOD concepts to virtualized geometry, neural implicit representations, radiance fields and 3D Gaussian splatting.

To understand the challenges and drawbacks of existing adaptive LOD systems especially in terms of the absence of a unified control based on perception and resource-awareness.

To suggest a conceptual approach and evaluation system for future adaptive LOD rendering systems.

Methodology

Research Design

The design of this study is qualitative and integrative literature-review. This paper is not based on a new user study or a rendering system, but is a synthesis of the

[1] Department of Computer Science, Air University, Islamabad, Pakistan.
[2] Department of Computer Science, Bahria University, Islamabad, Pakistan.

Copyright (c) 2024 Author(s)
This work is licensed under a Creative Commons Attribution 4.0 International License.

peer-reviewed papers from the conferences, papers in journals, survey literature and technical papers on LOD and real-time rendering. This design is suitable for the research question as it investigates the current status of adaptive LOD methods, their performance trade-offs and conceptual needs for future systems.

Source Selection

The criteria used to select sources were: relevance to real-time rendering, direct involvement with LOD or adaptive detail, methodological clarity and publication in a reputable scholarly or technical journal. Theoretical groundwork is laid by Clark (1976), Hoppe (1996), Garland and Heckbert (1997) and Luebke (2001). Virtualized Geometry, Neural rendering, Gaussian splatting are recent developments that are covered from 2021 to 2025, such as the works of Takikawa et al. (2021), Karis et al. (2021), Gu et al. (2023), Kerbl et al. (2023, 2024), Song et al. (2024), Haydel and Yuksel (2023) and Chen et al. (2025).

Analytical Framework

Five dimensions were considered, namely representation type, LOD granularity, adaptation criterion, transition mechanism, and reported performance focus, with each one being explored for each source. Polygonal meshes are distinguished by representation type as are implicit fields, radiance fields, and Gaussian splats. Object-level, cluster-level, block-level and per-ray adaptation are all distinguished by LOD granularity. The criteria for adaptation are distance, screen space error, occlusion, motion, hardware, and bandwidth criteria. There are several types of transition mechanisms: discrete, blending, interpolation, and refinement. Performance focus is comprised of frame rate, latency, memory, bandwidth, energy, image quality and temporal stability.

Proposed Conceptual Framework

To motivate this, a four-component adaptive LOD controller (scene hierarchy, cost model, perceptual importance model, and a budget-constrained selector) is proposed based on the synthesis. The scene hierarchy arranges geometry, Gaussians or neural blocks into a cluster of multiple detail levels. The cost model is a predictive calculation of the rendering cost of each candidate representation, based on primitive count, shaders complexity, memory traffic, and traversal depth.

[1] Department of Computer Science, Air University, Islamabad, Pakistan.
[2] Department of Computer Science, Bahria University, Islamabad, Pakistan.

Copyright (c) 2024 Author(s)
This work is licensed under a Creative Commons Attribution 4.0 International License.

The perceptual importance model sums the likely contribution of each cluster based on projected size, velocity, contrast, eccentricity and occlusion. The selector selects representations that provide the greatest perceptual gain, given a frame-time and memory-budget constraint.

The controller can be considered as solving an optimization problem with constraints. Define the selected detail level for scene element i as \(d_i\), the contribution of the element's perceptual quality as \(Q_i(d_i)\), and the expected cost of rendering the element as \(C_i(d_i)\).Define the detail level of the i-th scene element as \(d_i\), the contribution of the perceptual quality of the i-th scene element as \(Q_i(d_i)\), and the expected rendering cost of the i-th scene element as \(C_i(d_i)\). The goal is to achieve the greatest perceptual gain, Maximize $\Sigma Q_i(d_i)$ subject to $\Sigma C_i(d_i) \leq B$.

It makes explicit the trade-off that has to be dealt with by practical LOD systems, namely that detail should be assigned to where there is the greatest perceptual return for its computational cost. The framework is not empirical, but it is used as a guide for the future implementation and evaluation.

Results

Synthesis of Technique Characteristics.

It is seen from the literature that the process has continued to move from the static object-level LOD to the fine-grained resource-oriented adaptation. Classical polygonal simplification is based on geometric error metrics and generates precomputed approximations, while the hierarchical approach can be used during run time to choose among the scene partitions. Virtualized geometry also improves the decentralization of the choice by choosing small clusters based on screen-space error. The neural method and the Gaussian method also use similar concepts with non-mesh representations, in the sense that feature cells, volumetric blocks or Gaussian clusters are picked up depending on the view and the resources.

One thing that has been shown through many studies is that scalability is a requirement for hierarchical organization. Hierarchies narrow search space during

[1] Department of Computer Science, Air University, Islamabad, Pakistan.
[2] Department of Computer Science, Bahria University, Islamabad, Pakistan.
Copyright (c) 2024 Author(s)
This work is licensed under a Creative Commons Attribution 4.0 International License.

rendering, allow coarse to fine traversal, and allow for a compact representation of distant content. In explicitly using a hierarchy of 3D Gaussians, Kerbl et al. (2024) manage to retain quality in very large scenes while adjusting to available resources. In the same way, the HLOD systems use the octree-like structures to choose proper spatial partitions and prevent popping and cracking by using interpolation (CVMP, 2023).

A second common result is that the screen-space error is more appropriate than world-space distance strictly for LOD selection. The size, form, roughness or significance of objects at a given distance can be very different. The use of virtualized geometry systems like Nanite, on the other hand, relies on screen space density, which means that the more geometry occupies screen space, the more detailed and the more geometry that becomes small, the more it is simplified.For virtualized geometry systems like Nanite, however, it is based on screen space density, so that more geometry is detailed when it's occupying a large area of the screen and is simplified aggressively when it becomes small (Karis et al., 2021).

A third conclusion is that of the temporal stability. Sometimes abrupt LOD switching can create "popping" effects, especially when moving the camera. This artifact is minimized when using Continuous LOD, interpolation and progressive refinement. The NGLOD supports continuous LOD by signed-distance interpolation (Takikawa et al., 2021) and the LOD transitions are smooth using hierarchical Gaussian splatting (Kerbl et al., 2024). The mechanisms show that perceived quality is not only a function of the accuracy of the selected LOD, but also the stability of the transitions between LODs.

Performance Trade-Offs

The literature suggests that adaptive LOD can achieve better performance in realtime by processing fewer geometry, neural query or Gaussian primitives per frame. The amount and kind of the benefit however, depends on the rendering pipeline. With reduced primitive counts, vertex processing, rasterization and memory traffic will be reduced with rasterization. In the context of ray tracing, simplified geometry can decrease the number of traversal of the acceleration-

[1] Department of Computer Science, Air University, Islamabad, Pakistan.
[2] Department of Computer Science, Bahria University, Islamabad, Pakistan.
Copyright (c) 2024 Author(s)
This work is licensed under a Creative Commons Attribution 4.0 International License.

structure and the memory bandwidth, and Haydel and Yuksel (2023) have shown that their locally adaptive ray-tracing LOD structure can save over an order of magnitude in energy consumption and render time when computational resources are adequate. For neural rendering, the network query, volumetric samples, or the number of active scene blocks can all be decreased using LOD, as demonstrated by City-on-Web that adopts LOD and dynamic loading to lower memory usage (Song et al., 2024).

There are costs associated with these benefits. Multiple LODs need to be prepared in advance and this will take time and space. To store intermediate representations, hierarchies need memory. With runtime selection, there is an overhead of traversal and decision. Additional interpolation or blending may be necessary when making continuous transitions. The benefit of adaptive control is not always beneficial: When the scene is small, the object is simple, or the camera path is fixed, then the overhead of the adaptive control can outweigh the savings. The evidence also reviewed indicates that LOD is more consistent at improving scalability than at improving peak image quality. High detail representations are still required at close view distances. Adaptive LOD aims to avoid any processing of unimportant material, thus allowing more resources to be devoted to more important or moving areas. This is consistent with the perceptual principles that lead to LOD: that computation should be given in accordance with visual significance, not uniform complexity of a scene.

Discussion

System Design implications for rendering.

The synthesis implies that successful Adaptive LOD systems should be conceived as closed-loop resource managers and not simply as fixed LOD selection processes. A renderer should consider frame time, GPU usage, LOD pressure and streaming status, and adjust the thresholds or budget. For instance, if frame time is above the target, the system may decrease the detail of low saliency or unimportant clusters first, and then decrease the detail of foreground clusters or

[1] Department of Computer Science, Air University, Islamabad, Pakistan.
[2] Department of Computer Science, Bahria University, Islamabad, Pakistan.
Copyright (c) 2024 Author(s)
This work is licensed under a Creative Commons Attribution 4.0 International License.

high saliency content. If there is enough headroom, the system can interpolate the detail progressively.

This is especially crucial on diverse equipment. The compute power, memory bandwidth, thermal power, and display resolution requirements vary widely between desktop GPUs, mobile devices, and VR headsets as well as cloud-rendering clients. An adaptive controller can adjust the same scene representation to be usable on multiple platforms, instead of having to create separate, manually-tuned assets for each platform.

Perceptual Considerations

Distance is not the only factor involved in perceptual adaptation. People with eccentricities are more sensitive to their surroundings than those without.A more eccentric person is more sensitive to details, more sensitive to movement, more sensitive to contrast, and more sensitive to the task. Peripheral vision is less sharp than foveal vision and high-frequency information may be obscured by high-speed motion. These observations are in favor of policies that keep the detail in the fovea, in the case of high-contrast boundaries, and for objects moving relative to one another quickly. One should be careful of applying perceptual models, however: wrong predictions might lead to visible simplifications in certain areas where users look at the image, especially when interacting with the camera, such as in robot navigation.

Future systems will thus need to merge the computational metrics with user-focused validation. Peak signal-to-noise (PSNR) and structural similarity (SSIM) are objective metrics that quantify image difference, but are not sufficient to capture the effect of popping, shimmering, or task performance. When LOD changes have a material impact on the experience, studies should incorporate temporal-quality measures, user ratings, and measures based on tasks, like target detection or navigation accuracy.

Limitations

There are a number of limitations on this paper. It is a review-based synthesis (not new experimental evidence). Reported performances are based on the original publications and may be based on specific scenes, hardware, implementation

[1] Department of Computer Science, Air University, Islamabad, Pakistan.
[2] Department of Computer Science, Bahria University, Islamabad, Pakistan.
Copyright (c) 2024 Author(s)
This work is licensed under a Creative Commons Attribution 4.0 International License.

decisions, and evaluation conditions and are not intended to be used directly as comparisons. Secondly, the proposed framework is hypothetical and has not yet been applied and tested. Third, given the fast pace of changes in neural rendering, the literature surveyed here might not represent the latest advances in the field or the most up-to-date optimisations of techniques for specific hardware.
Future research should include the proposed controller in an open rendering engine and test against a set of controlled scene data sets. To assess frame time, memory bandwidth, geometric error, image quality, temporal stability and user perception, a rigorous study should measure outcomes as it varies the LOD policy, perceptual weighting and resource budgets. This evaluation would allow one to determine the usefulness of unified adaptive control over the conventional distance-based LOD for real-world rendering tasks.

Conclusion

One of the most crucial approaches for achieving optimal performance for real-time rendering is still adaptive level-of-detail (ALOD). LOD allows an interactive system to faithfully portray a scene that would be too time, memory or power consuming if portrayed in full detail. The works in the literature demonstrate a series of static polygonal simplification, hierarchical detail abstraction, view-dependent detail abstraction, virtual abstraction and neural abstraction.
The review reveals the key requirements of scalable real-time LOD, which are hierarchical organization, screen-space error metrics, smooth transitions. It also points to a large gap, namely that the current systems optimize the individual representation or local geometry decisions without having a unified, perceptually aware, resource-constrained control policy. This paper fills this gap by introducing the conceptual framework proposed here that views LOD selection as constrained optimization over perceptual benefit and rendering cost.
Future work should be focused on eliminating distance based substitution and building adaptive LOD controllers that take into account combined geometry and perceptual, motion, occlusion, streaming, and hardware constraints. These systems will be vital to the integration of dense meshes, volumetric effects, neural

[1] Department of Computer Science, Air University, Islamabad, Pakistan.
[2] Department of Computer Science, Bahria University, Islamabad, Pakistan.
Copyright (c) 2024 Author(s)
This work is licensed under a Creative Commons Attribution 4.0 International License.

radiance fields, and Gaussian-based representations in large-scale, dynamic virtual environments for interactive applications.

References

Chen, Y., et al. (2025). Virtualized 3D Gaussians: Flexible cluster-based level-of-detail system for real-time rendering of composed scenes. In SIGGRAPH Conference Papers '25. ACM. https://doi.org/10.1145/3721238.3730602

Clark, J. H. (1976). Geometric models for visible surface algorithms (hierarchical). Communications of the ACM, 19(10), 547–554. https://doi.org/10.1145/360349.360354

Garland, M., & Heckbert, P. S. (1997). Simplification of surfaces based on quadric error. In Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '97) (pp. 209–216). ACM. https://doi.org/10.1145/258734.258849

Gu, J., Jiang, M., Li, H., Lu, X., Zhu, G., Shah, S. A. A., Zhang, L., & Bennamoun, M. (2023). UE4-NeRF: Real-Time Rendering of a Large-Scale Scene using a Neural Radiance Field. V. M. Gulati, M. Ahmad, and M. N. S. H. Woodford in Advances in Neural Information Processing Systems 36 (NeurIPS 2023).

Haydel, J., & Yuksel, C. (2023). Locally-adaptive level-of-detail for hardware accelerated ray tracing. In SIGGRAPH Asia 2023 Conference Papers. ACM.

Hoppe, H. (1996). Progressive meshes. In Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '96) (pp. 99–108). ACM. https://doi.org/10.1145/237170.237216

Karis, B., Stubbe, R., & Wihlidal, G. (2021). Detailed exploration of virtualized geometry in Nanite. Advances in Real-Time Rendering in Games SIGGRAPH 2021 Courses. ACM.

Kerbl, B., Kopanas, G., Leimkühler, T., & Drettakis, G. (2023). Real-time rendering of a radiance field with 3D Gaussian splatting. ACM Transactions on Graphics, 42(4), 1–14. https://doi.org/10.1145/3592434

The authors would like to acknowledge the support of the European Union.The authors would like to thank the European Union for its support. (2024). Real-time

[1] Department of Computer Science, Air University, Islamabad, Pakistan.
[2] Department of Computer Science, Bahria University, Islamabad, Pakistan.

Copyright (c) 2024 Author(s)
This work is licensed under a Creative Commons Attribution 4.0 International License.

rendering of very large datasets using a hierarchical 3D Gaussian representation. ACM Transactions on Graphics, 43(4), Article 1. https://doi.org/10.1145/3666079

Luebke, D. (2001). A survey of polygonal simplification algorithms for the developer. IEEE Computer Graphics and Applications, 21(3), 24–35. https://doi.org/10.1109/38.920624

Luebke, D., Reddy, M., Cohen, J. D., Varshney, A., Watson, B., & Huebner, R. (2003). The degree of detail for the 3D graphics. Morgan Kaufmann.

Song, K., Zeng, X., Ren, C., & Zhang, J. (2024). City-on-Web: Real-time neural rendering on the Web for large-scale scenes. arXiv. https://arxiv.org/abs/2312.16457

Takikawa T., Litalien J., Yin K., Kreis K., Loop C., Nowrouzezahrai D., Jacobson A., McGuire M., Fidler S. (2021). Neural geometric level of detail: Implicit 3D shapes realtime rendered. In IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (pp., 12858–12867). IEEE.

Xu, L., Agrawal, V., Laney, W., Garcia, T., Bansal, A., Kim, C., Bulò, S. R., Porzi, L., Kontschieder, P., Božič, A., Lin, D., Zollhöfer, M., & Richardt, C. (2023). High fidelity virtualized walkable spaces: VR-NeRF. arXiv. https://arxiv.org/abs/2311.02542

[1] Department of Computer Science, Air University, Islamabad, Pakistan.
[2] Department of Computer Science, Bahria University, Islamabad, Pakistan.

Copyright (c) 2024 Author(s)
This work is licensed under a Creative Commons Attribution 4.0 International License.