Chapter 9 / Capítulo 9

SAP Superintelligence Series
ISSN: 3144-5624
ISBN (complete work): 978-9915-9851-0-7

Superintelligence Series Volume 5: Intelligent Visual Arts and Audiovisual Engineering: Generative AI, Computational Creativity and Superintelligence
ISBN (this volume): 978-9915-9966-5-3
DOI: 10.62486/978-9915-9966-3-9.Ch09
Pages: 109-120

Check for updates

© 2026 The authors. This is an open access book chapter distributed under the terms of the Creative Commons Attribution (CC BY) 4.0 License.

Neural radiance fields and 3D Gaussian splatting in audiovisual production

Campos de radiancia neuronal y salpicado gaussiano 3D en la producción audiovisual

Fiorella Luz De La Cruz Espinoza[1]
Norma Beatriz Ríos[2]
Carlos Jesús Canova Barrios[3,4]

1 Universidade Federal do Rio Grande do Sul. Brazil.
2 Universidad Hospital Italiano de Buenos Aires. Buenos Aires, Argentina.
3 Universidad Abierta Interamericana. Buenos Aires, Argentina.
4 Universidad de Ciencias Empresariales y Sociales. Buenos Aires, Argentina.

ABSTRACT

Neural rendering has introduced two dominant approaches to synthesizing photorealistic views of a scene, neural radiance fields and 3D Gaussian splatting, whose different trade-offs matter directly for audiovisual production. This narrative review synthesizes the two approaches across five axes, representation and training, rendering quality, rendering speed, editability, and dynamic-scene support, drawing on the literature published between 2018 and 2026. Neural radiance fields represent a scene as a continuous implicit function optimized by volume rendering, which yields high quality and strong anti-aliasing but is slow to train and render without baking or acceleration. Three-dimensional Gaussian splatting represents the same scene as an explicit set of anisotropic Gaussians rendered by rasterization, which delivers real-time frame rates with comparable quality but introduces artifacts and editing challenges of its own. The two approaches also differ in editability and in their treatment of moving scenes. The synthesis concludes that the choice between them is a production decision rather than a technical ranking, and that the field's open problems, temporal consistency, production-grade editability and standardized evaluation, are the likely routes by which neural rendering will enter routine audiovisual workflows.

Keywords: neural rendering; neural radiance fields; 3D Gaussian splatting; novel view synthesis; audiovisual production; real-time rendering

RESUMEN

El renderizado neuronal ha introducido dos enfoques dominantes para sintetizar vistas fotorrealistas de una escena, los campos de radiancia neuronal y el salpicado gaussiano 3D, cuyas distintas compensaciones importan directamente para la producción audiovisual. Esta revisión narrativa sintetiza ambos enfoques en cinco ejes, representación y entrenamiento, calidad de renderizado, velocidad, editabilidad y soporte de escenas dinámicas, a partir de la bibliografía publicada entre 2018 y 2026. Los campos de radiancia neuronal representan la escena como una función implícita continua optimizada mediante renderizado volumétrico, lo que ofrece alta calidad y buen antialiasado, pero es lento de entrenar y renderizar sin aceleración. El salpicado gaussiano 3D representa la misma escena como un conjunto explícito de gaussianas anisótropas renderizadas por rasterización, lo que logra velocidades en tiempo real con calidad comparable, aunque introduce sus propios artefactos y retos de edición. Los dos enfoques también difieren en editabilidad y en el tratamiento de escenas en movimiento. La síntesis concluye que la elección entre ambos es una decisión de producción, no una clasificación técnica, y que los problemas abiertos del campo, consistencia temporal, editabilidad de grado productivo y evaluación estandarizada, son las vías probables por las que el renderizado neuronal entrará en los flujos audiovisuales habituales.

Palabras clave: renderizado neuronal; campos de radiancia neuronal; salpicado gaussiano 3D; síntesis de vistas; producción audiovisual; renderizado en tiempo real

INTRODUCTION

Audiovisual production has always depended on the ability to render a three-dimensional scene as a two-dimensional image from an arbitrary camera, and for most of its history that ability has been bought with explicit geometry, meshes, materials and lights assembled by hand. Neural rendering changes the terms of that bargain. Instead of modeling a scene explicitly, a neural rendering method stores the scene inside a learned function or a learned set of primitives and synthesizes new views by querying it, which collapses the distance between capture and image and opens the way to free-viewpoint video, virtual production and post-capture relighting without a conventional three-dimensional pipeline.[1,2] The capacity to reconstruct a photorealistic scene from photographs and then move the camera through it is directly relevant to every stage of audiovisual work, from previsualization and virtual sets to visual effects, and the field has therefore turned toward these methods with unusual speed.(3,4)

Two approaches dominate this movement, and they represent opposite answers to the same question of how a scene should be stored. The first, the neural radiance field, encodes a scene as a continuous implicit function, typically a multilayer perceptron, that maps a position and a viewing direction to a color and a density, and it is optimized through differentiable volume rendering.[1] The second, three-dimensional Gaussian splatting, encodes the scene as an explicit cloud of anisotropic Gaussians, each carrying position, covariance, opacity and color, that are projected and rasterized in order.[2] The two approaches were introduced within three years of each other and now split the literature between them, with the Gaussian representation overtaking the radiance field as the default for real-time work after 2023.[3,5] Their contrast is informative precisely because they are competitors for the same production tasks, and because they fail in different, complementary ways.

The comparison matters for production because the two approaches sit on different sides of the quality-speed frontier that audiovisual work constantly navigates. A radiance field produces images of high fidelity and handles anti-aliasing and unbounded scenes cleanly, but its volume rendering queries the network hundreds of times per ray, which makes naive rendering far from real time.[1,6] Gaussian splatting reverses that relationship, rasterizing millions of primitives to reach real-time frame rates at full resolution while retaining comparable quality, at the cost of new artifacts such as popping and floaters.[2,7] For a supervisor choosing between a method that is slow but stable and one that is fast but noisy, the distinction is not academic, and neither the graphics literature nor the production literature has yet given practitioners a clear, comparative account of it.

The moment is also right for such an account, because the field has just passed through a reorientation that changed what the two methods are competing for. The radiance field founded the field in 2020 and carried it through its first expansion, but the Gaussian representation, introduced in 2023, overtook it as the default for real-time work with a speed that the literature itself remarks on, and the surveys that followed had to be written around two lineages rather than one.[3,4] A comparison made earlier would have been premature, because the Gaussian side of the ledger was still empty; a comparison made now can draw on three years of methods on both sides of the divide, and it arrives at a time when production houses are actively deciding which representation to invest in for virtual production, visual effects and free-viewpoint work.

Surveys of neural rendering have grown abundant, but they are organized around a single representation rather than around the contrast between the two. Comprehensive reviews trace the radiance-field lineage and its extensions, and parallel surveys trace Gaussian splatting and its applications, yet a reader who wants to know how the two actually compare for production must assemble the answer from separate accounts.(5,8-10) A smaller number of works confront the two directly, either in benchmark evaluations or in method papers that position one against the other, and those scattered comparisons are the raw material this review assembles.[3,4] The gap, then, is not a shortage of evidence but the absence of a synthesis that is organized by the contrast itself.

This review addresses that gap by comparing neural radiance fields and three-dimensional Gaussian splatting across five axes that correspond to the decisions a production actually makes: how the scene is represented and trained, how good the rendered image is, how fast it renders, how editable it is, and how well it handles moving scenes. The objective is to compare the two approaches as neural rendering methods for audiovisual production and to identify, from the comparison, the open problems that most limit their adoption in practice. Each axis is developed in the section that follows, and the five answers are integrated in the discussion into a single account of where each approach fits within a production workflow.

https://doi.org/10.62486/978-9915-9966-3-9.Ch09

METHODS

The review followed the SANRA guidance for narrative reviews, which means that the search and selection are described with transparency but without the formal screening counts of a systematic review.(11) The object of the synthesis is the primary literature on the two neural rendering approaches, together with the surveys and evaluation studies that position one against the other. The review is comparative by design, and its method is to hold a fixed set of production-relevant questions against each approach in turn and to report where the literature converges, where it diverges, and where it is silent.

Selection followed three criteria applied to title and abstract. A study was retained if it introduced, evaluated or surveyed a neural radiance field or Gaussian splatting method with direct relevance to one of the five axes; if it was published within the search window; and if it was written in English or Spanish. Works on neural rendering for unrelated tasks, such as medical imaging or autonomous navigation, were admitted only where they contributed a finding that transfers to audiovisual production, such as real-time performance or dynamic-scene handling.[12,13] The retained material was divided into a core of primary method papers that populate the synthesis and a supporting set of surveys and adjacent studies that supply context in the introduction and discussion.

The synthesis is organized around five contrast axes, each a question that a production team would ask. The first axis concerns representation and training, the second rendering quality, the third rendering speed, the fourth editability and control, and the fifth dynamic scenes and temporal consistency. For each axis the review reports what the two approaches offer, cites the primary studies that establish each capability, and notes the divergences in the literature where studies disagree about the magnitude of an advantage. Because the two approaches are measured on heterogeneous benchmarks with heterogeneous metrics, the synthesis is narrative rather than quantitative, and no pooled estimate is attempted; where a number is reported, it is attributed to the study that produced it.

A synthesis table (Table 1) summarizes the five axes with the representative studies that support each finding, and a second table (Table 2) sets the two approaches side by side on the production-relevant properties that recur across the axes.

Table 1. Synthesis of the five contrast axes.

Contrast axis	Neural radiance field	3D Gaussian splatting
Representation and training	Continuous field optimized by volume rendering	Explicit Gaussians rasterized in order
Rendering quality	High fidelity, clean anti-aliasing, fails under sparse input	Comparable fidelity, real time, prone to floaters
Rendering speed	Real time only with baking or compression	Real time natively at 1080p
Editability and control	Indirect editing via stylization, palettes or layers	Explicit primitives, easier but not mesh-grade
Dynamic scenes	Mature deformation and topology, human performance	Real-time dynamic playback, temporal consistency open

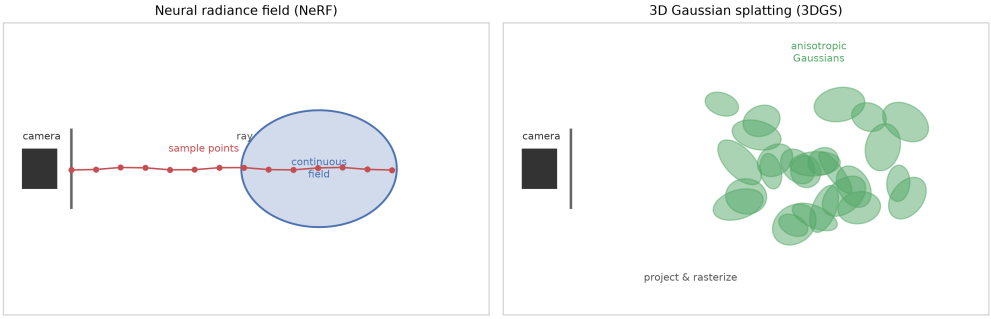
Table 2. Production-relevant properties of the two neural rendering approaches.

Property	Neural radiance field	3D Gaussian splatting
Scene representation	Continuous implicit function	Explicit set of anisotropic Gaussians
Training time	Hours, minutes with hash encoding	Tens of minutes
Rendering speed	Offline, real time only with baking	Real time at full resolution
Editability	Indirect, bespoke methods	Explicit primitives, partial editing
Dynamic support	Mature deformation and topology	Real-time playback, weaker temporal consistency
Primary failure	Sparse input, transient capture	Floaters, popping, complex scenes

A conceptual figure (Figure 1) illustrates the structural difference between the two representations, a continuous field sampled by ray marching and a discrete set of Gaussians projected by splatting, since that structural difference is the source of most of the axes that follow.

https://doi.org/10.62486/978-9915-9966-3-9.Ch09

Figure 1. The two neural rendering representations compared in this review.



Note: The left panel shows a radiance field sampled by ray marching, and the right panel shows a cloud of anisotropic Gaussians projected by splatting.

DEVELOPMENT

REPRESENTATION AND TRAINING

The two approaches differ first in what a scene is. A neural radiance field is a continuous function, a multilayer perceptron that takes a five-dimensional coordinate, three for position and two for viewing direction, and returns the volume density and view-dependent color at that point. (1) A view is synthesized by marching along each camera ray, sampling the function at many points, and compositing the samples by the volume rendering equation, and the network is optimized so that the rendered views reproduce the input photographs. This is a compact and flexible representation, but the optimization is per-scene and, in the original formulation, takes hours on a single scene, and the rendering step queries the network hundreds of times per ray, which is the root of both the quality and the cost of the method.

Gaussian splatting stores the same scene in the opposite way, as an explicit cloud of anisotropic three-dimensional Gaussians, each of which carries a position, a covariance describing its shape, an opacity and a set of spherical-harmonic coefficients describing its color, initialized from the sparse point cloud produced by structure from motion and then optimized against the photographs.(2) Rendering projects the Gaussians onto the image plane, sorts them, and composites them front to back by alpha blending, a rasterization step that runs on the graphics pipeline rather than through a network. The consequence is that training converges in tens of minutes rather than hours, and rendering reaches interactive rates without the acceleration tricks that radiance fields require.(2)

The training gap has been narrowed from both sides, and the narrowing itself is part of the comparison. On the radiance-field side, the multiresolution hash encoding and related feature-grid encodings replaced the coordinate network's raw input with a learned feature lookup, cutting per-scene optimization from hours to seconds or minutes while preserving most of the quality.(14) Baking approaches take the further step of precomputing a trained radiance field into a sparse grid that can be rendered without the network at all.(6) On the Gaussian side, the explicit representation arrived already fast, and the recent work has concentrated on scaling it to very large datasets through hierarchical and level-of-detail structures rather than on closing a training gap.[15,16] The axis therefore ends less in a clean victory than in a convergence: the two representations began far apart in cost and have been pulled toward the same middle by different mechanisms.

RENDERING QUALITY

On rendering quality the two approaches are close enough that the differences are matters of regime rather than of rank. The original radiance field demonstrated that a coordinate network could reproduce view-dependent appearance, including reflections and translucency, at a fidelity that was the point of the field's founding, and the anti-aliased variant of the method, which traces cones rather than rays, extended that fidelity to unbounded outdoor scenes.[1,17] Gaussian splatting reports quality on the same benchmark families that is comparable to the best radiance-field results while rendering at real-time rates, a combination that had previously been thought to require a quality trade-off.(2) The primary studies that evaluate both representations on shared test scenes consistently place them within a small margin of each other, with the ranking depending

https://doi.org/10.62486/978-9915-9966-3-9.Ch09

on the dataset and on the metric used, a point that matters more for benchmark reporting than for the production decision it is meant to inform.

The failures differ in kind, and a production reader needs the failures more than the aggregate scores. Radiance fields degrade most visibly under sparse input, where the geometry is under-constrained and the network produces floaters and holes, a weakness addressed by geometric regularization and by depth or RGB-D priors.[18,19] They also struggle with in-the-wild captures that contain transient occluders and moving objects, which transient modeling attempts to inpaint.[20] Gaussian splatting degrades most visibly in the other direction, not from under-constrained geometry but from the brittleness of its optimization and its sorting, which produces popping artifacts, floaters and view-dependent inconsistencies that dedicated methods target with sorted splatting, distractor removal and level-of-detail selection.(7,16,21) In short, radiance fields fail where the input is thin, and Gaussians fail where the scene is complicated, and neither failure is the same as the other.

A further quality distinction concerns the treatment of view-dependent appearance, which production work relies on more than benchmark suites reveal. The radiance field is built around a viewing direction as an explicit input, which is why reflections, translucency and specular highlights are among its strengths, and why anti-aliased variants extend those strengths to outdoor, unbounded scenes.[1,17] Gaussian splatting models view dependence through spherical-harmonic coefficients on each primitive, which is sufficient for many scenes but is the source of specific artifacts when reflections and transmission are prominent, and dedicated Gaussian formulations exist to model planar transmission and reflection precisely because the base representation handles them imperfectly.(22) This is not a ranking so much as a reminder that the two representations encode the same optical phenomena differently, and that a production that cares about glass, water or polished surfaces should weight the difference accordingly.

RENDERING SPEED

Rendering speed is the axis on which the two approaches are furthest apart and the one that has driven the field's reorientation. A naive radiance field is far from real time because every pixel requires the network to be evaluated hundreds of times along its ray, and early real-time results were bought by reducing resolution, by baking the network into a grid, or by compressing the representation into memory-efficient forms that fit a browser or a mobile device.(1,6,23) These are genuine real-time radiance fields, but they are the product of an extra engineering pass, and each acceleration trades away some of the flexibility of the original implicit representation.

Gaussian splatting is real time by construction. Because its primitives are projected and rasterized on the graphics pipeline rather than sampled through a network, the method reports real-time display at full high-definition resolution on consumer hardware, and this is the single property most responsible for its rapid adoption.(2) Subsequent work extends the real-time claim outward, to very large datasets through hierarchy, to extended-reality workloads through architecture co-design, and to planar transmission and reflection through dedicated Gaussian formulations.(15,22,24) The speed axis is thus the clearest asymmetry in the comparison: radiance fields can be made real time, but Gaussians are real time, and the distinction is felt in any interactive production context.

A second speed distinction cuts across both representations, between methods that train per scene and methods that generalize. Per-scene methods, which dominate both lineages, must re-optimize for each new capture, and that optimization time, whether minutes or hours, is a fixed entry cost.[1,2] Generalizable methods learn a mapping from images to a scene representation once and then synthesize views of unseen scenes in a forward pass, at the cost of some fidelity.(25-27) For production, the difference is between a method that is slow once and fast thereafter, and a method that is fast always but slightly less faithful, a distinction the axis preserves rather than collapses.

The acceleration techniques developed for radiance fields deserve separate mention because they illustrate how far the representation can be pushed before it stops being the representation. Real-time radiance fields for single-image portraits reach twenty-four frames per second on consumer hardware by predicting a triplane representation directly from a single photograph, and visual-hull acceleration reduces the cost of the volume rendering step by restricting sampling to the space a scene is known to occupy.[28,29] These methods matter for production because they target the two constraints that most limit on-set use, the need for a fast, per-shot result and the need to start from ordinary, unposed footage rather than from a

https://doi.org/10.62486/978-9915-9966-3-9.Ch09

calibrated multi-view rig. The radiance field is therefore not simply slow; it is slow in a way that a series of published accelerations has progressively, and only partially, dismantled.

EDITABILITY AND CONTROL

Editability is where both approaches reveal their immaturity, and where the contrast between implicit and explicit is most instructive. A radiance field is a network, and a network offers no direct handle on the objects it has learned, so editing must be achieved indirectly, through text-driven stylization that adjusts the field's appearance, through palette mappings that relight or recolor a frozen scene, or through layered decompositions that separate a dynamic scene into editable entities.(30-32) These are real editing capabilities, but they are each bespoke, and none of them gives a production artist the familiar operations of moving, deleting or reshaping an object that a mesh or a rig would provide.

Gaussian splatting is more promising on this axis because its primitives are explicit, and it inherits from point-based graphics the ability to deform, manipulate and even reconstruct a surface from its Gaussians in real time.[33,34] Yet the promise is only partially realized. The Gaussians are not objects, but millions of disconnected primitives, and editing them at the level a production expects, selecting a character, retiming a performance, re-lighting a shot, still requires custom machinery, since the primitives carry no semantic grouping and no articulation that an artist could grab. The axis therefore lands on a specific conclusion: the explicit representation is easier to edit than the implicit one, but neither has yet reached the editability of the conventional pipeline it aspires to replace.

DYNAMIC SCENES AND TEMPORAL CONSISTENCY

Moving scenes are the axis on which the audiovisual stakes are highest, because production is concerned with people and cameras in motion rather than with static objects. The radiance-field lineage addressed motion early and thoroughly. Non-rigid formulations model a scene as a canonical volume deformed over time, hyper-dimensional approaches add a time-varying topology to that canonical model, and streaming representations decompose a four-dimensional scene into static, deforming and new regions so that it can be played back as video.(35-37) Human performance is the clearest production target, and human-centric radiance fields demonstrate free-viewpoint playback of full-body motion from multi-view video at a fidelity aimed at film, with related work reconstructing human performance through neural meshes and generalizable human renderers.(38-40)

Gaussian splatting entered dynamic scenes after its static success, and it brought the same real-time property with it. Four-dimensional Gaussian splatting renders moving scenes in real time by making the Gaussians time-varying, and motion-aware variants mine optical flow to reduce the redundancy that dynamic scenes introduce.[41,42] Human rendering with Gaussians follows the same arc, from per-subject optimization to generalizable pixel-wise approaches that render a novel person from sparse multi-view capture without re-optimization.[43,44] The dynamic axis does not produce a clean winner: radiance fields offer the more mature treatment of deformation and topology, while Gaussians offer the real-time playback, and the literature has not yet combined both in a single method that production could adopt without reservation.

The axis closes on the problem that remains open for both, temporal consistency. A dynamic scene is judged not frame by frame but across time, and both representations are reported to flicker, pop and drift between frames under motion, a defect that is acknowledged in the dynamic literature and that no included method claims to have eliminated.[37,41] Temporal consistency is therefore the clearest single gap on the production path, and it is carried forward into the discussion as the problem the next generation of methods will most plausibly address, since a scene that renders beautifully in a still but flickers in playback is not yet a production asset.

DISCUSSION

The comparison across the five axes resolves into a single, stable picture: neural radiance fields and Gaussian splatting are not a better and a worse method, but two methods that trade the same two properties, fidelity and speed, in opposite directions, and that fail in different regimes. A radiance field buys fidelity and clean anti-aliasing with a slow, implicit representation

https://doi.org/10.62486/978-9915-9966-3-9.Ch09

that is costly to train, costly to render and difficult to edit, while Gaussian splatting buys real-time rendering and an explicit, manipulable representation at the cost of new artifacts and of an optimization that is brittle on complex scenes. This reading aligns with the surveys that describe the Gaussian representation as overtaking the radiance field after 2023 for real-time work, while acknowledging that the radiance field retains advantages in unbounded scenes and in the maturity of its dynamic treatment.[3,4] The review's contribution is to state the trade-off explicitly and to locate each production decision within it.

The quality axis deserves emphasis because it is the one on which a casual reader could misread the literature. The two approaches report comparable aggregate scores, but the review's reading is that they are not equally robust: radiance fields fail when the input is sparse, and Gaussians fail when the scene is complicated, and the surveys that report only aggregate metrics obscure this asymmetry.[5,10] The discrepancy with a purely benchmark-driven comparison is explained by what the benchmarks measure, which is average error on fixed scenes rather than failure behavior at the margins, and the review therefore treats the failure regimes as more informative for production than the mean scores. This is a case where the method of comparison, by contrast rather than by ranking, changes the conclusion.

The speed advantage of Gaussian splatting is the least contested finding in the review and the one with the clearest production consequence. Where a radiance field must be baked, compressed or accelerated to reach real time, Gaussian splatting reaches it natively, and the difference is felt in any context where an operator must interact with the scene, virtual production, previsualization and on-set review being the obvious examples.(2,6,23) The caveat, which the axis records, is that real-time is not the same as correct, and the artifacts that Gaussian methods accept in exchange for speed are exactly the defects that a broadcast-grade deliverable cannot tolerate. The practical implication is that the two methods may coexist in a single pipeline, with Gaussians for interaction and radiance fields, or a baked radiance field, for the final render.

The editability finding is the most sobering for the claim that neural rendering will soon replace conventional pipelines, and it should be stated without softening. Neither representation offers the operations that a production artist expects from a mesh, and while the explicit Gaussian representation is easier to edit than the implicit radiance field, both remain far from the editability of a rigged model, a point the editing literature confirms from its own survey of the field.(45) The review's reading is that editability, not rendering, is the true bottleneck for adoption, because a photorealistic render that cannot be art-directed is a novelty rather than a tool, and the production literature on artificial intelligence in film makes the same point from the practitioner's side.(46) The two approaches will enter routine production not when they render well, but when they can be directed.

The dynamic-scene axis connects the comparison to the specific concerns of audiovisual work, where the subject is usually a person in motion. The radiance-field lineage has the more mature treatment of deformation, topology and human performance, with dedicated human-centric methods aimed explicitly at film, while Gaussian splatting contributes real-time dynamic playback that the radiance-field methods cannot match without heavy acceleration.(35,38,41) The axis identifies a gap rather than a winner, and the gap is temporal consistency, the tendency of both representations to flicker and drift between frames, which no included method claims to have eliminated and which is the most direct technical obstacle between the current state of the field and a broadcast deliverable.[37,42] This gap is the review's central recommendation for future work.

The comparison also carries a caution that the technical literature rarely states but that the audiovisual context forces forward. Neural rendering does not only synthesize views, it synthesizes views of people, and the same capacity that produces free-viewpoint performance can produce fabricated performance, a concern the documentary community has already confronted through the practice of deepfakes.(47) The ethical dimension is not an afterthought to the comparison but a property of the technology that production adopts, and a review that ignores it would understate what is actually at stake when a photorealistic, directable scene representation becomes routine. The two approaches are not distinguished by this concern, but the concern attaches to both, and it reinforces the review's emphasis on editability and control as the property that separates a legitimate production tool from a fabrication instrument.

The axes are not abstract, and two application domains make their practical force concrete. In cultural-heritage documentation, radiance fields have been evaluated across heritage datasets captured with smartphone video and tourist photography, where the finding is that accelerated variants of the method reach the quality a documentation workflow needs, a use that values fidelity over interactivity and therefore favors the radiance-field side of the trade-off.(48) In virtual-reality

https://doi.org/10.62486/978-9915-9966-3-9.Ch09

applications, by contrast, non-photorealistic rendering is used precisely as a feedback mechanism that keeps a user immersed, a reminder that photorealism is not always the production goal and that the speed and interactivity of a representation can matter more than its fidelity in a given application.(49) These domains illustrate the review's central claim in reverse: the correct choice of representation is decided by the production axis that a task weights most, not by a universal ranking of the two methods.

The review also sits in relation to the surveys that already exist, and that relation should be stated plainly. Overviews of the radiance field trace its technology and applications, and parallel surveys of Gaussian splatting do the same for the explicit representation, but each treats one lineage, and a practitioner comparing the two must read across them.(5,8,9,50) This review is distinguished by holding the contrast fixed, so that the same question is asked of both methods in the same section, which is what turns a pair of adjacent literatures into a single comparison. The comparison itself is the contribution, and the five axes are the instrument that makes it possible.

For a production pipeline, the comparison resolves into a division of labor rather than a choice of one method over the other. Gaussian splatting's native real-time rendering suits the interactive stages, previsualization, on-set review and virtual production, where an operator must move through the scene and the cost of an artifact is low because the image is provisional.[2,24] A radiance field, or a baked radiance field, suits the final render, where the camera is fixed, the image must be clean, and the slow, stable fidelity of the implicit representation is worth the wait.[6,17] The two methods are therefore not substitutes but complements, and the practical implication is that a pipeline that adopts both, Gaussians for interaction and fields for delivery, gains more than one that commits to either.

The comparison also locates neural rendering within the larger advance of artificial intelligence into audiovisual production, and that larger frame is part of what gives the technical trade-off its significance beyond the computer-graphics venues in which the methods were first reported. Neural rendering is not an isolated technique but one point on a curve that runs from automated editing and synthesis to the directable, photorealistic reconstruction of a scene from photographs, and the production literature already reads artificial intelligence as a force moving through pre-production, production and post-production as a whole.(46) What distinguishes neural rendering on that curve is that it concerns the image itself rather than the tasks around the image, which is why its open problems, temporal consistency, editability and evaluation, are also the points at which a general-purpose AI capacity becomes a dependable production tool. The choice between the two representations is therefore a specific instance of a broader decision that production is making about which AI capabilities to adopt and which to treat as provisional, and the review's axes are the concrete terms in which that decision is made.

The agenda that follows from the review is specific rather than general:

- Temporal consistency in dynamic scenes is the single highest-value technical problem, because it is the one that separates current neural rendering from broadcast use, and it is shared by both approaches; a method that holds a moving scene stable across frames would do more for adoption than any incremental improvement in mean fidelity.
- Production-grade editability, meaning selection, deformation, retiming and relighting of semantic entities rather than of millions of primitives or of a frozen network, is the property most likely to determine whether either approach is adopted, and it is best pursued by attaching semantic grouping to the explicit Gaussian representation, which already offers the primitives that grouping could bind.
- The field needs a shared benchmark that reports failure regimes as well as mean error, so that the quality comparison the review could only make qualitatively can be made quantitatively, and so that the complementary failure modes identified here can be measured rather than asserted.
- A fourth item follows from the capture side of the comparison, standardizing the input from which production scenes are reconstructed, since the sparse and unposed footage that production actually provides is not the calibrated multi-view input that both representations assume.

These elements are the routes by which the comparison in this review would close, and they are offered as concrete, testable directions for research rather than as the generic call for more study.

https://doi.org/10.62486/978-9915-9966-3-9.Ch09

CONCLUSIONS

Neural radiance fields and three-dimensional Gaussian splatting were compared across the five axes that matter to audiovisual production, and the comparison shows that the two approaches trade fidelity against speed in opposite directions while failing in different regimes. The radiance field offers a continuous implicit representation with strong quality, clean anti-aliasing and a mature treatment of dynamic scenes, at the cost of slow training, slow naive rendering and difficult editing. Gaussian splatting offers an explicit representation with native real-time rendering and easier manipulation, at the cost of new artifacts and a brittle optimization on complex scenes. On dynamic scenes the two approaches are distinguished less by a winner than by a shared gap, temporal consistency, which remains unresolved for both. The review concludes that the choice between them is a production decision governed by the specific axis that a given project weights most heavily, and that the field's route into routine use runs through three problems in particular: temporal consistency, production-grade editability and an evaluation practice that reports failure as carefully as it reports mean error.

https://doi.org/10.62486/978-9915-9966-3-9.Ch09

REFERENCES

Mildenhall B, Srinivasan PP, Tancik M, Barron JT, Ramamoorthi R, Ng R. NeRF. Commun ACM. 2022;65(1):99-106. doi: 10.1145/3503250
Kerbl B, Kopanas G, Leimkuehler T, Drettakis G. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. ACM Trans Graph. 2023;42(4):1-14. doi: 10.1145/3592433
Gao K, Gao Y, He H, Lu D, Xu L, Li J. Neural Radiance Fields in 3D Vision: A Comprehensive Review. Comp Visual Med. 2026;12(4):851-905. doi: 10.26599/cvm.2026.9450535
Fei B, Xu J, Zhang R, Zhou Q, Yang W, He Y. 3D Gaussian Splatting as a New Era: A Survey. IEEE Trans Visual Comput Graphics. 2025;31(8):4429-4449. doi: 10.1109/tvcg.2024.3397828
Chen G, Wang W. A Survey on 3D Gaussian Splatting. ACM ComputSurv. 2026;58(12):1-39. doi: 10.1145/3807511
Hedman P, Srinivasan PP, Mildenhall B, Reiser C, Barron JT, Debevec P. Baking Neural Radiance Fields for Real-Time View Synthesis. IEEE Trans Pattern Anal Mach Intell. 2025;47(5):3310-3321. doi: 10.1109/tpami.2024.3381001
Radl L, Steiner M, Parger M, Weinrauch A, Kerbl B, Steinberger M. StopThePop: Sorted Gaussian Splatting for View-Consistent Real-time Rendering. ACM Trans Graph. 2024;43(4):1-17. doi: 10.1145/3658187
Liao Y, Di Y, Zhou H, Zhu K, Lu M, Duan Q, et al. A Survey on Neural Radiance Fields. ACM ComputSurv. 2026;58(2):1-33. doi: 10.1145/3758085
Xiao W, Chierchia R, Santa Cruz R, Li X, Ahmedt Aristizabal D, Salvado O, et al. Neural Radiance Fields for the Real World: A Survey. ACM ComputSurv. 2026;58(15):1-36. doi: 10.1145/3837758
Bao Y, Ding T, Huo J, Liu Y, Li Y, Li W, et al. 3D Gaussian Splatting: Survey, Technologies, Challenges, and Opportunities. IEEE Trans Circuits Syst Video Technol. 2025;35(7):6832-6852. doi: 10.1109/tcsvt.2025.3538684
Chen W, Li Z, Guo J, Zheng C, Tian S. Trends and Techniques in 3D Reconstruction and Rendering: A Survey with Emphasis on Gaussian Splatting. Sensors. 2025;25(12):3626. doi: 10.3390/s25123626
Wang G, Pan L, Peng S, Liu S, Xu C, Miao Y, et al. NeRFs in robotics: A survey. The International Journal of Robotics Research. 2026;45(6):858-892. doi: 10.1177/02783649251374246
Lei X, Wang M, Zhou W, Li H. GaussNav: Gaussian Splatting for Visual Navigation. IEEE Trans Pattern Anal Mach Intell. 2025;47(5):4108-4121. doi: 10.1109/tpami.2025.3538496
Müller T, Evans A, Schied C, Keller A. Instant neural graphics primitives with a multiresolution hash encoding. ACM Trans Graph. 2022;41(4):1-15. doi: 10.1145/3528223.3530127
Kerbl B, Meuleman A, Kopanas G, Wimmer M, Lanvin A, Drettakis G. A Hierarchical 3D Gaussian Representation for Real-Time Rendering of Very Large Datasets. ACM Trans Graph. 2024;43(4):1-15. doi: 10.1145/3658160
Ren K, Jiang L, Lu T, Yu M, Xu L, Ni Z, et al. Octree-GS: Towards Consistent Real-time Rendering with LOD-Structured 3D Gaussians. IEEE Trans Pattern Anal Mach Intell. 2025;PP:1-15. doi: 10.1109/tpami.2025.3568201
Barron JT, Mildenhall B, Verbin D, Srinivasan PP, Hedman P. Mip-NeRF 360: Unbounded Anti-Aliased Neural Radiance Fields. 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 2022:5460-5469. doi: 10.1109/cvpr52688.2022.00539
Niemeyer M, Barron JT, Mildenhall B, Sajjadi MSM, Geiger A, Radwan N. RegNeRF: Regularizing Neural Radiance Fields for View Synthesis from Sparse Inputs. 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 2022:5470-5480. doi: 10.1109/cvpr52688.2022.00540
Yuan YJ, Lai YK, Huang YH, Kobbelt L, Gao L. Neural Radiance Fields From Sparse RGB-D Images for High-Quality View Synthesis. IEEE Trans Pattern Anal Mach Intell. 2023;45(7):8713-8728. doi: 10.1109/tpami.2022.3232502
Wang S, Xu H, Li Y, Chen J, Tan G. IE-NeRF: Exploring transient mask inpainting to enhance neural radiance fields in the wild. Neurocomputing. 2025;618:129112. doi: 10.1016/j.neucom.2024.129112
Sabour S, Goli L, Kopanas G, Matthews M, Lagun D, Guibas L, et al. SpotLessSplats: Ignoring Distractors in 3D Gaussian Splatting. ACM Trans Graph. 2025;44(2):1-11. doi: 10.1145/3727143
Liu Y, Ye K, Shao T, Zhou K. TR-Gaussians: High-Fidelity Real-Time Rendering of Planar Transmission and Reflection With 3D Gaussian Splatting. IEEE Trans Visual Comput Graphics. 2026;32(7):5882-5894. doi: 10.1109/tvcg.2026.3675416
Reiser C, Szeliski R, Verbin D, Srinivasan P, Mildenhall B, Geiger A, et al. MERF: Memory-Efficient Radiance Fields for Real-time View Synthesis in Unbounded Scenes. ACM Trans Graph. 2023;42(4):1-12. doi: 10.1145/3592426
Chang K, Gan Y, Zhu W, Wang K, Chen Z, Cheng Y, et al. ScaleGS: Closing the Gap between Real-time 3D Gaussian Splatting and Real-time XR Rendering. ACM Trans Archit Code Optim. 2025;22(4):1-26. doi: 10.1145/3774425
Yu A, Ye V, Tancik M, Kanazawa A. pixelNeRF: Neural Radiance Fields from One or Few Images. 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 2021:4576-4585. doi: 10.1109/cvpr46437.2021.00455
Nguyen-Ha P, Huynh L, Rahtu E, Matas J, Heikkilä J. Cascaded and Generalizable Neural Radiance Fields for Fast View Synthesis. IEEE Trans Pattern Anal Mach Intell. 2024;46(5):2758-2769. doi: 10.1109/tpami.2023.3335311
Jiang L, Mao Y, Xu L, Lu T, Ren K, Jin Y, et al. AnySplat: Feed-forward 3D Gaussian Splatting from Unconstrained Views. ACM Trans Graph. 2025;44(6):1-16. doi: 10.1145/3763326
Trevithick A, Chan M, Stengel M, Chan E, Liu C, Yu Z, et al. Real-Time Radiance Fields for Single-Image Portrait View Synthesis. ACM Trans Graph. 2023;42(4):1-15. doi: 10.1145/3592460
Marí R. Accelerating NeRF with the Visual Hull. Image Processing On Line. 2024;14:217-231. doi: 10.5201/ipol.2024.553
Wang C, Jiang R, Chai M, He M, Chen D, Liao J. NeRF-Art: Text-Driven Neural Radiance Fields Stylization. IEEE Trans Visual Comput Graphics. 2024;30(8):4983-4996. doi: 10.1109/tvcg.2023.3283400
Wu Q, Tan J, Xu K. PaletteNeRF: palette-based color editing for NeRFs. Communications in Information and

Systems. 2023;23(4):447-475. doi: 10.4310/cis.2023.v23.n4.a4
Zhang J, Liu X, Ye X, Zhao F, Zhang Y, Wu M, et al. Editable free-viewpoint video using a layered neural representation. ACM Trans Graph. 2021;40(4):1-18. doi: 10.1145/3450626.3459756
Gao L, Yang J, Zhang BT, Sun JM, Yuan YJ, Fu H, et al. Real-time Large-scale Deformation of Gaussian Splatting. ACM Trans Graph. 2024;43(6):1-17. doi: 10.1145/3687756
Lyu X, Sun YT, Huang YH, Wu X, Yang Z, Chen Y, et al. 3DGSR: Implicit Surface Reconstruction with 3D Gaussian Splatting. ACM Trans Graph. 2024;43(6):1-12. doi: 10.1145/3687952
Tretschk E, Tewari A, Golyanik V, Zollhofer M, Lassner C, Theobalt C. Non-Rigid Neural Radiance Fields: Reconstruction and Novel View Synthesis of a Dynamic Scene From Monocular Video. 2021 IEEE/CVF International Conference on Computer Vision (ICCV). 2021:12939-12950. doi: 10.1109/iccv48922.2021.01272
Park K, Sinha U, Hedman P, Barron JT, Bouaziz S, Goldman DB, et al. HyperNeRF. ACM Trans Graph. 2021;40(6):1-12. doi: 10.1145/3478513.3480487
Song L, Chen A, Li Z, Chen Z, Chen L, Yuan J, et al. NeRFPlayer: A Streamable Dynamic Scene Representation with Decomposed Neural Radiance Fields. IEEE Trans Visual Comput Graphics. 2023;29(5):2732-2742. doi: 10.1109/tvcg.2023.3247082
Işık M, Rünz M, Georgopoulos M, Khakhulin T, Starck J, Agapito L, et al. HumanRF: High-Fidelity Neural Radiance Fields for Humans in Motion. ACM Trans Graph. 2023;42(4):1-12. doi: 10.1145/3592415
Zhao F, Jiang Y, Yao K, Zhang J, Wang L, Dai H, et al. Human Performance Modeling and Rendering via Neural Animated Mesh. ACM Trans Graph. 2022;41(6):1-17. doi: 10.1145/3550454.3555451
Gao X, Yang J, Kim J, Peng S, Liu Z, Tong X. MPS-NeRF: Generalizable 3D Human Rendering From Multiview Images. IEEE Trans Pattern Anal Mach Intell. 2025;47(8):6110-6121. doi: 10.1109/tpami.2022.3205910
Wu G, Yi T, Fang J, Xie L, Zhang X, Wei W, et al. 4D Gaussian Splatting for Real-Time Dynamic Scene Rendering. 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 2024:20310-20320. doi: 10.1109/cvpr52733.2024.01920
Guo Z, Zhou W, Li L, Wang M, Li H. Motion-Aware 3D Gaussian Splatting for Efficient Dynamic Scene Reconstruction. IEEE Trans Circuits Syst Video Technol. 2025;35(4):3119-3133. doi: 10.1109/tcsvt.2024.3502257
Zheng S, Zhou B, Shao R, Liu B, Zhang S, Nie L, et al. GPS-Gaussian: Generalizable Pixel-Wise 3D Gaussian Splatting for Real-Time Human Novel View Synthesis. 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 2024:19680-19690. doi: 10.1109/cvpr52733.2024.01861
Zhou B, Zheng S, Tu H, Shao R, Liu B, Zhang S, et al. GPS-Gaussian+: Generalizable Pixel-wise 3D Gaussian Splatting for Real-Time Human-Scene Rendering from Sparse Views. IEEE Trans Pattern Anal Mach Intell. 2025;PP:1-16. doi: 10.1109/tpami.2025.3561248
Hubert A, Elghazaly G, Frank R. Editing implicit and explicit representations of radiance fields: a survey. Machine Vision and Applications. 2025;36(6):119. doi: 10.1007/s00138-025-01742-4
Davaasambuu T. Film Production Artificial Intelligence Use. ADR. 2025;13(03):210-220. doi: 10.4236/adr.2025.133016
Lees D. Deepfakes in documentary film production: images of deception in the representation of the real. Studies in Documentary Film. 2024;18(2):108-129. doi: 10.1080/17503280.2023.2284680
Mazzacca G, Karami A, Rigon S, Farella EM, Trybala P, Remondino F. NERF FOR HERITAGE 3D RECONSTRUCTION. Int Arch Photogramm Remote Sens Spatial Inf Sci. 2023;XLVIII-M-2-2023:1051-1058. doi: 10.5194/isprs-archives-xlviii-m-2-2023-1051-2023
Quintana D, Rodríguez A, Sbert M, Silva Y, Rufino E, Boada I. Non-photorealistic rendering as a feedback strategy in virtual reality for rehabilitation. Virtual Reality. 2024;28(1):60. doi: 10.1007/s10055-024-00954-9
Chen T, Yang Q, Chen Y. Overview of NeRF Technology and Applications. J Comput-Aided Design Comput Graphics. 2025;37(1):51-74. doi: 10.3724/sp.j.1089.2023-00759

https://doi.org/10.62486/978-9915-9966-3-9.Ch09