EUROGRAPHICS 2026/ R. Chaine and A. Vaxman *Tutorial*

Fast Explicit 3D Reconstructions and How To Use Them

B. Kerbl[1], L. Franke[2], F. Hahlbohm[3], M. Steinberger[4], and A. Tagliasacchi[5]

[1]University of Copenhagen, Denmark
[2]INRIA, France
[3]TU Braunschweig, Germany
[4]Graz University of Technology, Austria
[5]Simon Fraser University, Canada

Abstract

Creating high-fidelity digital twins from photographs has become increasingly practical thanks to recent advances in explicit 3D scene representations. Among these, 3D Gaussian Splatting (3DGS) stands out for its combination of high visual quality and extremely fast rendering. However, deploying such models in real applications reveals several challenges. A growing body of research addresses these limitations from multiple angles. Extensions to 3DGS improve portability, scalability, and stability; specialized techniques target flicker, popping, and aliasing; neural point-based methods introduce principled encodings that avoid several drawbacks of Gaussian primitives; and alternative explicit representations——including triangles, stochastic splats, and novel volumetric structures—offer different trade-offs for varied use cases. This full-day tutorial provides a clear, consolidated overview of these fast explicit 3D representations. Attendees will learn how the major families of methods work, what constraints they address, and how to choose the right representation for specific interactive reconstruction or exploration tasks. Through conceptual explanations and end-to-end workflows, participants will gain practical guidance for using these techniques in their own research or pipelines, along with an understanding of their current limitations and open challenges.

Keywords: Real-time graphics, radiance fields, splatting, point-based graphics

CCS Concepts
• ***Computing methodologies*** → ***Machine learning; Rendering; Parallel computing methodologies;***

1. Introduction

The ability to create a high-fidelity 3D copy of a real-world scene—a digital twin—from a set of photographs poses an exciting, well-established challenge for visual computing. The ability to produce such explorable 3D representations enables a plethora of practical use cases, e.g., in surveying, city planning, robotics, e-commerce, and the entertainment industry. Following the explosive success of Neural Radiance Fields (NeRF) [MST*21] for Novel-View Synthesis (NVS) and 3D reconstruction tasks, a significant body of follow-up work sought to improve on its interpretability and speed. These efforts have given rise to a range of fast 3D representations that abandon NeRF's *implicit* nature, and instead opt for *explicit* primitives to encode spatially varying information. 3D Gaussian Splatting [KKLD23] (3DGS) is a notable example of these explicit techniques, enabling high quality and extremely fast performance, for both optimization and rendering. While the method has proven to be highly versatile, it faces numerous challenges. Due to its large memory consumption, 3DGS models are not fit for deployment on low-end or portable devices, prohibiting, e.g., portable use on mobile phones or head-mounted displays for virtual reality (VR). In addition, optimization requires a significant compute effort, exhibiting very little control over the process. Last but not least, the custom primitive representation demands a custom rendering pipeline with support for volumetric blending. While such properties can often be tolerated for research purposes, they impede the real-world usability of obtained digital twins. The answers for addressing these challenges come in many different flavors: several methods extend 3DGS directly to make it more controllable, scalable, and portable [NSW24, PKK*24, TRS*25, MGK*24, KMK*24, KRS*24, SPH*25]. Specific attention is given to the avoidance of inherent artifacts under motion, such as popping and aliasing [CKD*25, SKR*25, RSP*24, HFW*25, TKdlT24, SGK*25]. Other point-based techniques explore principled projections, hierarchical encodings, and neural representations, which can inherently prevent several of 3DGS's drawbacks [ASK*20, RFS22, HKT*23, FRF*23, FRFS24, HFK*25, HFO*25, KHS*25, FFS25]. Yet another direction explores changing the reconstruction primitives altogether, splatting more traditional graphics primitives (such as triangles), stochastic samples, or fast traversal of expressive volumetric primitives (such as Voronoi cells) [GRYT25, HVD*25, KVK*25, TGF*25, GL24a].

Proceedings published by Eurographics - The European Association for Computer Graphics.

DOI: 10.2312/egt.20261003

The purpose of this tutorial is to provide a comprehensive overview over the developments, advances, and *specific use cases* that these varied techniques offer. Confronted with a given set of constraints, the audience of this tutorial will be thoroughly informed which of the available, fast 3D representations best suits their needs for interactive scene exploration and reconstruction, and how to deploy them for their own research or workflow. Furthermore, they will be recognizant of the challenges that these representations face, and what constitutes the practical limits w.r.t. cutting-edge research in this domain. This raises awareness for effective, yet lesser-known techniques, enabling participants to make the best possible choice for their use case.

2. Syllabus

The syllabus is planned for a full-day tutorial, and covers four major sections in total, with slots accounting for introduction, break, and questions from the audience.

- Welcome & Speaker Introduction
- Coffee Break
- Part 1—3DGS: Running, Compressing, Converting
- Coffee Break
- Part 2—Battling Artifacts: Removing Flicker, Pop, and Ghosts
- Lunch Break
- Part 3—Neural Points: Small Points, Crisp Details
- Coffee Break
- Part 4—Beyond Points: Triangles, Foams, and More
- Q & A

We note that the exact contents of the following syllabus sections is subject to change, since additional work in this domain has been submitted by various presenters, and the manuscripts are currently under review. We omit these methods from the description in order to ensure compliance with double-blind anonymous reviewing.

2.1. 3DGS: Running, Compressing, Converting

3D Gaussian Splatting has become a well-established method and tool for many visual computing workflows. Understanding the concepts it is built on (elliptic weighted average (EWA) splatting, point-based rendering, parallel software rendering) and its challenges (file size, portability, artifacts) is key for understanding subsequent methods. Further, this section is intended to equip novices to novel-view synthesis and image-based reconstruction with a basic understanding of the process. In addition to an introduction to 3DGS' background, as well as a demonstration of the original workflow (code setup + execution), this portion also introduces variants based on the original 3DGS codebase for compressing, accelerating, budgeting, and converting 3DGS models. Specifically, it will cover the specific requirements and examples for using

1. 3DGS [KKLD23], the original codebase and user interface
2. Reduced3DGS [PKK*24], which reduces the file size of obtained scenes by more than 30×
3. Taming3DGS [MGK*24], which enables prioritization and the strict budgeting of run-time memory and model size
4. Hierarchical3DGS [KMK*24], which extends 3DGS from living room to city scale by using hierarchical level-of-detail and on-demand streaming
5. SuGaR [GL24b], which provides a sophisticated solution for obtaining deformable meshes from 3D Gaussian Splatting models.

Beyond the basics and motivation for the original work, the main insights of this tutorial portion will enable users to tweak 3D Gaussian Splatting to their needs in constrained use cases, while maintaining the seminal Gaussian Splat representation.

2.2. Battling Artifacts: Removing Flicker, Pop, and Ghosts

Much like its predecessors, 3DGS in the wild suffers from a variety of issues, such as changing lighting and ephemeral details. While being an inherent artifact of the monocular capturing setup, it can be successfully addressed by cutting-edge techniques [SGK*25]. However, 3DGS also introduces its own set of aliasing challenges, which justify a detailed analysis [CKD*25]. While these artifacts are often not perceptible in desktop setups, extreme camera settings and motion (e.g., in VR applications) significantly exacerbate their disruptiveness. This section both explores the underlying reasons for various sources of aliasing in 3DGS and further presents rendering and training methods to prevent them at runtime. Starting from an initial discussion of principled behavior in volumetric rendering (and how 3DGS deviates from it), this portion of the tutorial will introduce techniques to combat them:

1. 3DGS artifacts and their impact on rendering [CKD*25]
2. SpotlessSplats [SGK*25], a technique for avoiding ghosting and residual artifacts of ephemeral objects in the capture
3. StopThePop [RSP*24], a fundamental solution to the approximate visibility algorithm that 3DGS employs
4. HTGS [HFW*25], a perspective-correct rendering approach for 3D Gaussians accelerated through hybrid transparency
5. AAA-Gaussians [SKR*25], a general revised rendering framework to address aliasing in 3DGS
6. VRSplat [TRS*25], a fast and robust solution that tackles visible aliasing in VR settings
7. Splatshop [SPH*25], an optimized toolbox for editing 3DGS models to fix reconstruction artifacts in post

At the end of this session, the audience will understand the design choices made in 3DGS, as well as the impact of shortcuts in the original algorithm that lead to occasional aliasing. Furthermore, they will be equipped to deploy rendering solutions that circumvent these problems.

2.3. Neural Points: Small Points, Crisp Details

We further introduce the principles of neural point rendering (NPR) and its progression towards modern 3DGS. Neural point rendering works by exploiting fast point sample rendering instead of larger splats, and solves the hole filling problem using screen-space filtering instead of world-space kernels (such as 3D Gaussians). This shifts the rendering complexity from rasterization to a neural reconstruction step. We begin with key NPR methods, such as NPBG [ASK*20] and ADOP [RFS22] pioneering this concept, and follow with techniques reducing the neural reconstruction step's overhead [FRFS24]. Because NPR and 3DGS both rely on point-based geometric proxies, the course covers methods for obtaining and refining point sets: dense point clouds from multi-view stereo or LiDAR, heuristic densification strategies, and VET [FRF*23] for error-driven densification during optimization. We also review INPC [HFK*25, HFO*25] and its dynamic extension D-NPC [KHS*25], which avoid reliance on explicit point clouds by learning implicit point structures. We conclude with hybrid approaches that combine neural point rendering with 3DGS. These are especially relevant to VR, where the tradeoff between primitive count, resolution, and rendering cost leads to a natural performance sweet spot [FFS25].

2.4. Beyond Points: Triangles, Foams, and More

While point-based representations have proven their ability for faithful and fast reconstruction, they are often not the ideal medium: for one, their lack of structure with Delta-like distribution spikes does not guarantee full coverage of the 3D reconstruction space. As a consequence, details can be entirely missed in reconstruction. On the other hand, the fundamental primitive in real-time 3D computer graphics is not the point, but the triangle; hardware vendors and engine developers have tailored real-time rendering designs, exploiting the long-standing convention of triangle meshes for efficient visibility resolution and rasterization.

This section of the tutorial focuses on emerging alternative representations with notable impact on usability. Radiant Foam [GRYT25] serves as an example for optimization on the entirety of a scene's volume, while maintaining highly efficient rendering. Triangle Splatting [HVD*25], on the other hand, directly yields triangle primitives; the resulting files can be immediately deployed as assets in 3D game engines without requiring custom rendering pipelines. Combining mesh extraction from Gaussian Splats (e.g., using SuGaR) with thin shells around detected object boundaries yields yet another representation—Gaussian Frosting [GL24a]—that is particularly useful for manipulation and animation while preserving fine-grained details, including support in the Blender3D authoring software. Finally, opacity fields [RWD*25] are specifically tailored for fast, high-fidelity mesh extraction.

At the end of this session, the audience will be thoroughly informed about explicit representations that omit point-based sampling, their strengths, model use cases, and promising opportunities for additional experiments and exploration.

3. Target Audience

The topics of this talk are fit for both beginners and advanced practitioners from academia and industry. Each section features an introduction for the individual techniques, providing theoretical backgrounds on the presented methods, as well as open challenges and opportunities. Furthermore, selected techniques will be presented with a full workflow from input to output and sample applications, enabling non-professionals and practitioners to easily understand the range of possible use cases and why individual techniques are tailor-made for particular challenges and constraints. After the initial introduction and discussion of 3D Gaussian Splatting, even novice listeners should be sufficiently prepared to follow the modifications that each subsequent technique introduces.

4. Speakers

The proposed speakers for this tutorial have contributed substantially to the space of high-performance 3D representations, radiance fields, and visual computing overall. Each speaker was strongly involved in the development of one or multiple such techniques, and thus possesses highly valuable insights into the challenges and opportunities of the presented representations. Their research accounts for a significant portion of the syllabus, which also includes important related work by unaffiliated contributors.

Bernhard Kerbl holds a PhD from Graz University of Technology and is a Tenure-track Assistant Professor at the University of Copenhagen. His research on high-performance processing yielded several state-of-the-art, GPU-based solutions for point-based rendering, including the seminal 3D Gaussian Splatting paper.

Linus Franke is a postdoctoral researcher working at Inria Sophia Antipolis in the GraphDeco group of George Drettakis. He holds a PhD from the Friedrich-Alexander-Universität Erlangen-Nürnberg supervised by Marc Stamminger. His research interests lie at the intersection of computer graphics, computer vision, and machine learning, with particular focus on point-based scene reconstruction and rendering.

Florian Hahlbohm is a junior researcher and PhD candidate at TU Braunschweig in the computer graphics group, led by Marcus Magnor. Florian's work focuses on advancing fast and precise point-based methods for image-based 3D reconstruction and view synthesis, achieving a balance between computational efficiency and visual quality.

Markus Steinberger is a full professor at Graz University of Technology (TU Graz). His research interests and his contributions to the domain of high-performance computer graphics and anti-aliased 3D Gaussian Splatting are reflected by the numerous awards won by his papers, including ACM CHI, IEEE Infovis, Eurographics, ACM NPAR, EG/ACM HPG, and IEEE HPEC best paper and honorable mentions.

Andrea Tagliasacchi is an Associate Professor at Simon Fraser University, where he holds the Visual Computing Research Chair. His work has been recognized with several distinctions for Best Paper and Best Student Paper Awards. Andrea's research focuses on challenges in 3D visual perception, a field that bridges computer vision, computer graphics, and machine learning.

References

[ASK*20] Aliev K.-A., Sevastopolsky A., Kolos M., Ulyanov D., Lempitsky V.: Neural point-based graphics. In *European Conference on Computer Vision (ECCV)* (Cham, 2020), Springer, Springer International Publishing, pp. 696–712. 1, 2

[CKD*25] Celarek A., Kopanas G., Drettakis G., Wimmer M., Kerbl B.: Does 3D Gaussian Splatting Need Accurate Volumetric Rendering? *Computer Graphics Forum* (2025). `doi:10.1111/cgf.70032`. 1, 2

[FFS25] Franke L., Fink L., Stamminger M.: VR-Splatting: Foveated Radiance Field Rendering via 3D Gaussian Splatting and Neural Points. *Proceedings of the ACM on Computer Graphics and Interactive Techniques 8*, 1 (May 2025). `doi:10.1145/3728302`. 1, 3

[FRF*23] Franke L., Rückert D., Fink L., Innmann M., Stamminger M.: VET: Visual Error Tomography for Point Cloud Completion and High-Quality Neural Rendering. In *SIGGRAPH Asia 2023 Conference Papers* (New York, NY, USA, 2023), SA '23, Association for Computing Machinery. `doi:10.1145/3610548.3618212`. 1, 3

[FRFS24] Franke L., Rückert D., Fink L., Stamminger M.: TRIPS: Trilinear Point Splatting for Real-Time Radiance Field Rendering. *Computer Graphics Forum 43*, 2 (2024), e15012. 1, 2

[GL24a] Guédon A., Lepetit V.: Gaussian frosting: Editable complex radiance fields with real-time rendering. *ECCV* (2024). 1, 3

[GL24b] Guédon A., Lepetit V.: Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering. *CVPR* (2024). 2

[GRYT25] Govindarajan S., Rebain D., Yi K. M., Tagliasacchi A.: Radiant foam: Real-time differentiable ray tracing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (October 2025), pp. 4135–4145. 1, 3

[HFK*25] Hahlbohm F., Franke L., Kappel M., Castillo S., Eisemann M., Stamminger M., Magnor M.: INPC: Implicit neural point clouds for radiance field rendering. In *International Conference on 3D Vision* (2025), pp. 168–178. `doi:10.1109/3DV66043.2025.00021`. 1, 3

[HFO*25] Hahlbohm F., Franke L., Overkämping L., Wespe P., Castillo S., Eisemann M., Magnor M.: A bag of tricks for efficient implicit neural point clouds. In *Vision, Modeling, and Visualization* (2025). `doi:10.2312/vmv.20251229`. 1, 3

[HFW*25] Hahlbohm F., Friederichs F., Weyrich T., Franke L., Kappel M., Castillo S., Stamminger M., Eisemann M., Magnor M.: Efficient perspective-correct 3d gaussian splatting using hybrid transparency. *Computer Graphics Forum 44*, 2 (2025). URL: `https://fhahlbohm.github.io/htgs/`, `doi:10.1111/cgf.70014`. 1, 2

[HKT*23] Hahlbohm F., Kappel M., Tauscher J.-P., Eisemann M., Magnor M.: Plenopticpoints: Rasterizing neural feature points for high-quality novel view synthesis. In *Proc. Vision, Modeling and Visualization (VMV)* (Sep 2023), Grosch T., Guthe M., (Eds.), Eurographics, pp. 53–61. `doi:10.2312/vmv.20231226`. 1

[HVD*25] Held J., Vandeghen R., Deliege A., Hamdi A., Cioppa A., Giancola S., Vedaldi A., Ghanem B., Tagliasacchi A., Van Droogenbroeck M.: Triangle splatting for real-time radiance field rendering. *arXiv* (2025). 1, 3

[KHS*25] Kappel M., Hahlbohm F., Scholz T., Castillo S., Theobalt C., Eisemann M., Golyanik V., Magnor M.: D-npc: Dynamic neural point clouds for non-rigid view synthesis from monocular video. *Computer Graphics Forum 44*, 2 (2025), e70038. URL: `https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.70038`, `arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.70038`, `doi:https://doi.org/10.1111/cgf.70038`. 1, 3

[KKLD23] Kerbl B., Kopanas G., Leimkühler T., Drettakis G.: 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics 42*, 4 (July 2023). URL: `https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/`. 1, 2

[KMK*24] 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 Transactions on Graphics 43*, 4 (July 2024). URL: `https://repo-sam.inria.fr/fungraph/hierarchical-3d-gaussians/`. 1, 2

[KRS*24] Kheradmand S., Rebain D., Sharma G., Sun W., Tseng Y.-C., Isack H., Kar A., Tagliasacchi A., Yi K. M.: 3d gaussian splatting as markov chain monte carlo. In *Proceedings of the 38th International Conference on Neural Information Processing Systems* (Red Hook, NY, USA, 2024), NIPS '24, Curran Associates Inc. 1

[KVK*25] Kheradmand S., Vicini D., Kopanas G., Lagun D., Yi K. M., Matthews M., Tagliasacchi A.: Stochasticsplats: Stochastic rasterization for sorting-free 3d gaussian splatting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (2025). 1

[MGK*24] Mallick S. S., Goel R., Kerbl B., Steinberger M., Carrasco F. V., De La Torre F.: Taming 3dgs: High-quality radiance fields with limited resources. In *SIGGRAPH Asia 2024 Conference Papers* (New York, NY, USA, 2024), SA '24, Association for Computing Machinery. URL: `https://doi.org/10.1145/3680528.3687694`, `doi:10.1145/3680528.3687694`. 1, 2

[MST*21] Mildenhall B., Srinivasan P. P., Tancik M., Barron J. T., Ramamoorthi R., Ng R.: Nerf: representing scenes as neural radiance fields for view synthesis. *Commun. ACM 65*, 1 (Dec. 2021), 99–106. URL: `https://doi.org/10.1145/3503250`, `doi:10.1145/3503250`. 1

[NSW24] Niedermayr S., Stumpfegger J., Westermann R.: Compressed 3d gaussian splatting for accelerated novel view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2024), pp. 10349–10358. 1

[PKK*24] Papantonakis P., Kopanas G., Kerbl B., Lanvin A., Drettakis G.: Reducing the memory footprint of 3d gaussian splatting. *Proc. ACM Comput. Graph. Interact. Tech. 7*, 1 (May 2024). URL: `https://doi.org/10.1145/3651282`, `doi:10.1145/3651282`. 1, 2

[RFS22] Rückert D., Franke L., Stamminger M.: ADOP: Approximate Differentiable One-Pixel Point Rendering. *ACM Transactions on Graphics (TOG) 41*, 4 (July 2022). `doi:10.1145/3528223.3530122`. 1, 2

[RSP*24] Radl L., Steiner M., Parger M., Weinrauch A., Kerbl B., Steinberger M.: StopThePop: Sorted Gaussian Splatting for View-Consistent Real-time Rendering. *ACM Transactions on Graphics 4*, 43 (2024). 1, 2

[RWD*25] Radl L., Windisch F., Deixelberger T., Hladky J., Steiner M., Schmalstieg D., Steinberger M.: SOF: Sorted Opacity Fields for Fast Unbounded Surface Rconstruction. In *SIGGRAPH Asia Conference Proceedings* (2025). 3

[SGK*25] Sabour S., Goli L., Kopanas G., Matthews M., Lagun D., Guibas L., Jacobson A., Fleet D., Tagliasacchi A.: Spotlesssplats: Ignoring distractors in 3d gaussian splatting. *ACM Trans. Graph. 44*, 2 (Apr. 2025). URL: `https://doi.org/10.1145/3727143`, `doi:10.1145/3727143`. 1, 2

[SKR*25] Steiner M., Köhler T., Radl L., Windisch F., Schmalstieg D., Steinberger M.: AAA-Gaussians: Anti-Aliased and Artifact-Free 3D Gaussian Rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (2025). 1, 2

[SPH*25] Schütz M., Peters C., Hahlbohm F., Eisemann E., Magnor M., Wimmer M.: Splatshop: Efficiently editing large gaussian splat models. *Computer Graphics Forum (Proc. HPG) 44*, 8 (2025). URL: `https://publications.graphics.tudelft.nl/papers/822`, `doi:10.1111/cgf.70214`. 1, 2

[TGF*25] Taktasheva M., Goli L., Fiorini A., Li Z., Rebain D., Tagliasacchi A.: 3d gaussian flats: Hybrid 2d/3d photometric scene reconstruction. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems* (2025). URL: `https://openreview.net/forum?id=uVxQEIgXfL`. 1

[TKdlT24] Tu X., Kerbl B., de la Torre F.: Fast and robust 3d gaussian splatting for virtual reality. In *SIGGRAPH Asia 2024 Posters* (New York, NY, USA, 2024), SA '24, Association for Computing Machinery. URL: `https://doi.org/10.1145/3681756.3697947`, `doi:10.1145/3681756.3697947`. 1

[TRS*25] Tu X., Radl L., Steiner M., Steinberger M., Kerbl B., de la Torre F.: Vrsplat: Fast and robust gaussian splatting for virtual reality. *Proc. ACM Comput. Graph. Interact. Tech. 8*, 1 (May 2025). URL: `https://doi.org/10.1145/3728311`, `doi:10.1145/3728311`. 1, 2

Proceedings published by Eurographics - The European Association for Computer Graphics.