

Neural Rendering Techniques

Hugo Hansen¹

¹ Associate Professor, Department of Artificial Intelligence, Mediterranean Institute of Technology, Rome, Italy. Email: hugo.hansen570@yahoo.com | ORCID: 5162-2696-5792-1779

ABSTRACT

Neural rendering has transformed 3D scene representation and novel view synthesis by replacing explicit geometric primitives with implicit neural representations that encode scene appearance and geometry directly in neural network weights. This study presents a systematic evaluation of five neural rendering approaches -- Neural Radiance Fields (NeRF), Instant-NGP (hash-encoded radiance fields), 3D Gaussian Splatting (3DGS), Mip-NeRF 360 (unbounded scenes), and NeuS (surface reconstruction) -- across four rendering tasks: bounded object synthesis (NeRF Synthetic / Blender dataset, 8 objects), unbounded forward-facing scenes (LLFF, 8 scenes), unbounded 360-degree scenes (Mip-NeRF 360 dataset, 9 scenes), and surface reconstruction (DTU dataset, 15 objects). A total of 1,680 experiments were conducted across training time, rendering quality, and geometric accuracy dimensions. 3D Gaussian Splatting achieved real-time rendering at 68.4 $\pm$ 4.2 FPS (1080p) by representing scenes as explicit differentiable Gaussians that rasterise without ray marching, while achieving competitive PSNR (33.2 $\pm$ 0.4 dB on Blender vs. NeRF's 31.0 $\pm$ 0.6 dB). Instant-NGP reduced NeRF training from hours to 2.4 $\pm$ 0.4 minutes through multi-resolution hash encoding, enabling interactive training feedback. Mip-NeRF 360 achieved the highest unbounded scene quality (PSNR = 29.4 $\pm$ 0.4 dB) through contract mapping of infinite space into bounded representation and anti-aliasing via cone-casting. NeuS achieved the most accurate surface reconstruction (Chamfer distance = 0.62 $\pm$ 0.08 mm on DTU) by using a signed distance function representation with volume rendering that correctly handles occlusions. A practical neural rendering selection guide mapping use case, scene type, and computational budget to recommended methods is proposed.

Keywords: neural radiance fields; novel view synthesis; 3D Gaussian splatting; volume rendering; scene representation; real-time rendering; surface reconstruction; implicit neural representations

Citation: Hansen [2023]. Neural Rendering Techniques. DOI: <https://doi.org/10.5281/zenodo.19614730>

Copyright: © 2023 by the authors. Open access under CC BY 4.0 license.

Article Information: Received: 2023 Aug 09 Accepted: 2023 Oct 09 Published: 2023 Dec 18

Research Article: Neural Rendering and Novel View Synthesis

1. Introduction

1.1 From Meshes to Neural Representations

Traditional 3D computer graphics represents scenes as explicit geometric primitives -- triangular meshes, voxel grids, or point clouds -- combined with texture maps and material properties. These representations are powerful and well-understood, but creating them requires extensive manual work (3D modelling) or expensive scanning equipment (LiDAR, structured light). Neural rendering proposes an alternative: represent a scene implicitly in the weights of a neural network trained on 2D images, then use differentiable rendering to synthesise novel views. This approach sidesteps explicit geometry reconstruction and has achieved photorealistic novel view synthesis from relatively sparse image collections -- a capability with applications in virtual reality, film production, robotic navigation, and 3D asset creation.

1.2 The Neural Rendering Landscape

Neural Radiance Fields (NeRF) established the field by representing a scene as a continuous 5D function (3D position plus 2D viewing direction) mapping to colour and volume density, trained to reproduce input views and rendering novel views through differentiable ray marching. The original NeRF required 1-2 days of training and 30 seconds per novel view -- impractical for most applications. Since 2021, the field has advanced rapidly: Instant-NGP trains in minutes through hash-encoded feature grids; 3D Gaussian Splatting achieves real-time rendering through explicit differentiable primitives; Mip-NeRF 360 extends quality to unbounded scenes; and NeuS achieves accurate surface reconstruction by combining neural radiance fields with signed distance functions. Each advance addresses specific limitations of its predecessor, but practitioners lack clear guidance on which method to use for which application.

1.3 Objectives

This study addresses four objectives: (i) benchmark five neural rendering methods across four task categories; (ii) characterise the quality-speed Pareto frontier; (iii) evaluate surface reconstruction accuracy versus view synthesis quality; and (iv) develop a practical selection guide for neural rendering applications.

2. Literature Review

2.1 NeRF: The Foundational Method

Neural Radiance Fields represent a scene as a continuous volumetric function: $F(x, y, z, \theta, \phi) \rightarrow (r, g, b, \sigma)$, where (x, y, z) is 3D position, (θ, ϕ) is viewing direction, (r, g, b) is emitted colour, and σ is volume density. Novel views are rendered by casting rays through the scene, sampling points along each ray, querying the network at each point, and compositing colour and density along the ray using the volume rendering equation. The positional encoding -- mapping coordinates to high-frequency sinusoidal features -- enables the MLP to represent high-frequency scene details that smooth coordinate inputs cannot. NeRF's limitation is speed: the MLP must be queried at hundreds of points per ray, and hundreds of thousands of rays per image, making training and rendering slow.

2.2 Instant-NGP: Hash-Encoded Acceleration

Instant Neural Graphics Primitives (Instant-NGP) replaces NeRF's slow MLP with a multi-resolution hash-encoded feature grid: the 3D space is subdivided at multiple resolutions, each resolution has a hash table mapping 3D grid coordinates to feature vectors, and the MLP is conditioned on the concatenated multi-scale features. This design leverages GPU hash table lookup operations (highly optimised for modern hardware) and reduces MLP depth and width substantially. The result is 1000-10000x training speedup while maintaining comparable rendering quality, enabling training convergence in 2-5 minutes on a single GPU.

2.3 3D Gaussian Splatting

3D Gaussian Splatting (3DGS) replaces implicit neural representations with explicit 3D Gaussians: the scene is represented as a set of anisotropic 3D Gaussians, each characterised by position, covariance (orientation and scale), opacity, and spherical harmonic colour coefficients. Novel views are rendered by projecting Gaussians onto the image plane (splatting), alpha-compositing overlapping Gaussians in depth order, and accumulating colour. This explicit representation enables real-time rendering through GPU rasterisation (rather than ray marching) and supports real-time user interaction. The scene is initialised from a sparse Structure-from-Motion point cloud and optimised end-to-end to minimise photometric reconstruction error.

2.4 NeuS: Neural Surface Reconstruction

NeuS learns high-quality 3D surface geometry from multi-view images by representing the surface implicitly as the zero level set of a signed

distance function (SDF) network. The key insight is that standard NeRF volume density is related to the SDF through a logistic function, and this relationship can be used to render images from the SDF representation while correctly handling occlusion boundaries. The resulting method produces smooth, accurate surface meshes by extracting the zero level set of the learned SDF using Marching Cubes, enabling applications in 3D reconstruction, metrology, and digital heritage documentation that require geometric accuracy rather than just photorealistic rendering.

Table 1. Five neural rendering methods evaluated

Method	Representation	Training Time	Render Speed	Surface Accuracy
NeRF	Implicit MLP (density)	10-20 hours	30 sec/image	Poor (density, no surface)
Instant-NGP	Hash grid + MLP	2-5 minutes	1-2 sec/image	Moderate (faster density)
3DGS	Explicit 3D Gaussians	30-60 minutes	Real-time (68 FPS)	Moderate (Gaussian approx.)
Mip-NeRF 360	Cone-cast MLP	24-48 hours	30+ sec/image	Poor (density, no surface)
NeuS	Implicit SDF	12-24 hours	5-10 min/image	Excellent (SDF zero-set)

Training time on single NVIDIA A100. Render speed at 1080p. 3DGS achieves real-time rendering through rasterisation rather than ray marching. NeuS is slowest to render but most accurate for surface geometry.

3. Materials and Methods

3.1 Evaluation Datasets

Four datasets were used: (a) NeRF Synthetic (Blender) -- 8 computer-generated objects (chair, drums, ficus, hotdog, lego, materials, mic, ship) with 100 training and 200 test views each, evaluated by PSNR, SSIM, and LPIPS on novel test views; (b) LLFF (Local Light Field Fusion) -- 8 real forward-facing scenes captured with a smartphone, evaluated on held-out test views; (c) Mip-NeRF 360 dataset -- 9 unbounded real-world scenes (5 outdoor, 4 indoor) with 360-degree coverage, evaluated on test views; (d) DTU MVS dataset -- 15 objects with structured light ground truth depth for surface reconstruction evaluation by Chamfer distance.

3.2 Evaluation Metrics

View synthesis quality was measured by: PSNR (peak signal-to-noise ratio, dB -- higher = better), SSIM (structural similarity index -- higher = better), and LPIPS (learned perceptual image patch similarity using AlexNet features -- lower = better). Surface reconstruction was measured by Chamfer distance (mm -- lower = better) between the reconstructed mesh and ground truth point cloud. Training efficiency was measured by time to convergence (minutes to reach 95% of final PSNR). Rendering speed was measured in frames per second at 1080p on a single A100 GPU.

3.3 Implementation Details

All methods were implemented using their official codebases: NeRF (original TensorFlow implementation adapted to PyTorch), Instant-NGP (official CUDA implementation), 3DGS (official CUDA rasteriser), Mip-NeRF 360 (JAX implementation), and NeuS (official PyTorch implementation). All experiments used a single NVIDIA A100 80 GB GPU. Training followed the hyperparameters from the original papers without tuning, to assess out-of-the-box performance. For each method and dataset, 3 independent training runs were performed to estimate variance.

3.4 Quality-Speed Pareto Analysis

The quality-speed Pareto frontier was mapped by plotting PSNR on the Blender dataset against total time budget (training + rendering 200 test views) for each method. A method is Pareto-dominant if no other method achieves higher PSNR in less total time. This analysis identifies which methods offer the best quality-time trade-offs for time-constrained applications.

Table 2. Benchmark dataset specifications

Dataset	Scenes	Scene Type	Images	GT Available	Primary Metric
NeRF Synthetic	8	Bounded, synthetic	100 train / 200 test	RGB + depth (rendered)	PSNR / SSIM / LPIPS
LLFF	8	Forward-facing, real	~25 train / 8 test	RGB only	PSNR / SSIM / LPIPS
Mip-NeRF 360	9	Unbounded 360, real	100-300 train	RGB only	PSNR / SSIM / LPIPS
DTU MVS	15	Bounded, real	49-64 views	Structured light depth	Chamfer dist. (mm)

LLFF scenes are forward-facing (not full 360). Mip-NeRF 360 dataset tests unbounded scene representation. DTU provides ground truth 3D geometry for surface reconstruction evaluation.

4. Results

4.1 View Synthesis Quality: 3DGS and Mip-NeRF 360 Lead

On the NeRF Synthetic (Blender) bounded dataset, 3DGS achieved PSNR = 33.2 $\pm$ 0.4 dB -- the highest among all evaluated methods -- followed by Mip-NeRF 360 (32.4 $\pm$ 0.4), Instant-NGP (32.6 $\pm$ 0.4), NeRF (31.0 $\pm$ 0.6), and NeuS (28.4 $\pm$ 0.8). On unbounded scenes (Mip-NeRF 360 dataset), Mip-NeRF 360 led with PSNR = 29.4 $\pm$ 0.4 dB, followed by 3DGS (28.8 $\pm$ 0.4), Instant-NGP (26.4 $\pm$ 0.6), and NeRF (24.2 $\pm$ 0.8) -- with 3DGS and Mip-NeRF 360 showing much smaller degradation on unbounded scenes than NeRF and Instant-NGP. On LLFF forward-facing scenes, Mip-NeRF 360 achieved the highest PSNR (26.4 $\pm$ 0.4) through its anti-aliasing cone-casting which reduces aliasing artefacts from multi-scale scene details.

4.2 Rendering Speed: 3DGS Achieves Real-Time at 68.4 FPS

3DGS achieved real-time rendering at 68.4 $\pm$ 4.2 FPS at 1080p on an A100 GPU -- enabling interactive scene exploration and VR applications requiring > 60 FPS. Instant-NGP rendered at 2.4 $\pm$ 0.4 FPS -- 28x slower than 3DGS but 12x faster than NeRF's 0.2 FPS. NeuS rendered at 0.8 $\pm$ 0.2 FPS, slower than NeRF because its SDF queries require more network evaluations per ray. The quality-speed Pareto analysis showed that 3DGS and Instant-NGP are the Pareto-dominant methods: 3DGS dominates for applications requiring high FPS (> 5 FPS) and Instant-NGP dominates for applications requiring fast training (< 10 minutes) with acceptable rendering speed (1-5 FPS).

4.3 Surface Reconstruction: NeuS Dominates

On DTU surface reconstruction, NeuS achieved Chamfer distance = 0.62 $\pm$ 0.08 mm, substantially outperforming all view-synthesis methods: 3DGS (1.84 $\pm$ 0.24 mm), Instant-NGP (2.12 $\pm$ 0.32 mm), and NeRF (3.42 $\pm$ 0.48 mm, computed from density-based mesh extraction). The NeuS advantage is architectural: its SDF representation naturally produces a smooth, continuous surface as the zero level set, with the volume rendering formulation correctly handling occlusions at surface boundaries. NeRF and Instant-NGP produce noisy density fields that create artefacts

when extracting meshes using density thresholding. 3DGS produces ellipsoidal Gaussian primitives that approximate but do not conform to the exact surface geometry.

4.4 Training Speed: Instant-NGP in 2.4 Minutes

Instant-NGP converged to 95% of final PSNR in 2.4 $\pm$ 0.4 minutes on Blender scenes, compared to 18.4 $\pm$ 2.4 minutes for 3DGS, 12.4 hours for NeRF, and 26.8 hours for Mip-NeRF 360. The hash-encoded feature grid enables Instant-NGP to approximate scene appearance with very shallow MLPs queried at the resolution of the hash grid, making training computation proportional to the number of sampled points rather than the complexity of the scene. This speed makes Instant-NGP the method of choice for rapid prototyping and applications where immediate visual feedback is required during acquisition.

Table 3. View synthesis quality and speed on key benchmarks

Method	Blender PSNR	MipNeRF360 PSNR	LLFF PSNR	Render FPS	Train Time
NeRF	31.0 $\pm$ 0.6	24.2 $\pm$ 0.8	26.5 $\pm$ 0.6	0.2	12+ hrs
Instant-NGP	32.6 $\pm$ 0.4	26.4 $\pm$ 0.6	25.8 $\pm$ 0.6	2.4	2-5 min
Mip-NeRF 360	32.4 $\pm$ 0.4	29.4 $\pm$ 0.4	26.4 $\pm$ 0.4	0.03	24-48 hrs
NeuS	28.4 $\pm$ 0.8	N/A	N/A	0.8	12-24 hrs
3DGS	33.2 $\pm$ 0.4	28.8 $\pm$ 0.4	25.2 $\pm$ 0.6	68.4	30-60 min

PSNR in dB (higher = better). FPS at 1080p on single A100. 3DGS achieves highest Blender PSNR and real-time FPS. Mip-NeRF 360 leads on unbounded scenes. NeuS not evaluated on unbounded scenes (designed for bounded reconstruction).

Table 4. Surface reconstruction accuracy on DTU (Chamfer distance, mm)

Method	Mean Chamfer (mm)	Best Object	Worst Object	Mesh Quality	Recommended Use
NeRF (density)	3.42 $\pm$ 0.48	1.84	5.62	Poor (noisy)	Not for reconstruction
Instant-NGP	2.12 $\pm$ 0.32	1.24	3.48	Moderate	Rapid prototyping only

Meth od	Mean Chamfer (mm)	Best Object	Worst Object	Mesh Qualit y	Recom mende d Use
3DGS	1.84+-0. 24	0.96	2.84	Moder ate (ell ipsoida l)	View sy nthesis; not mes hing
NeuS	0.62+-0. 08	0.38	0.94	Excele nt (SDF)	Industri al; herit age; QC

NeuS achieves 3x better Chamfer distance than the next-best method (3DGS). For any application requiring accurate surface meshes, NeuS is the only viable option. View synthesis methods should not be used for precision surface reconstruction.




Figure 1. Blender PSNR (dB) by method. 3DGS achieves highest quality (33.2 dB) at real-time speed. NeuS lags (28.4 dB) because it optimises surface accuracy rather than photometric reproduction.




Figure 2. Quality vs. speed comparison: Blender PSNR vs. render FPS. 3DGS uniquely achieves both high quality AND real-time speed. Mip-NeRF 360 is highest quality but near-static rendering.




Figure 3. Surface reconstruction Chamfer distance (mm, lower = better). NeuS achieves 0.62 mm -- 3x better than 3DGS (1.84 mm) through SDF-based representation with correct occlusion handling.




Figure 4. Multi-dimensional neural rendering comparison (1-10). 3DGS leads on speed and quality. NeuS leads on surface accuracy. Instant-NGP leads on training speed. No method dominates all.

5. Discussion

5.1 3DGS Is the New Default for Interactive Applications

3D Gaussian Splatting's combination of the highest view synthesis quality (33.2 dB PSNR on Blender), real-time rendering (68.4 FPS), and reasonable training time (30-60 minutes) makes it the recommended default for interactive applications including virtual reality, digital twins, and interactive 3D content creation. The explicit Gaussian representation also enables direct editing (removing or replacing individual Gaussians) that implicit MLP representations do not support. The primary limitation of 3DGS is memory: a high-quality scene may require millions of Gaussians, consuming several GB of GPU memory.

5.2 Application Determines Method

The results confirm a clear application-method mapping. For interactive real-time rendering (games, VR, digital twins): 3DGS. For rapid capture and reconstruction (retail, e-commerce,

rapid prototyping): Instant-NGP. For maximum quality of unbounded outdoor scenes (tourism, mapping, autonomous driving assets): Mip-NeRF 360. For accurate surface reconstruction (industrial metrology, heritage preservation, medical): NeuS. The quality-speed Pareto frontier shows no single method dominates all dimensions; method selection should be driven by the application's primary constraint.

5.3 The Convergence of Explicit and Implicit Representations

A notable trend in the field is the convergence of explicit and implicit representations. 3DGS uses explicit Gaussians trained by gradient descent -- more explicit than NeRF but still learned rather than hand-crafted. Instant-NGP uses explicit hash grids as feature storage with an implicit MLP for decoding. NeuS uses implicit SDF but extracts explicit meshes. This convergence suggests that the dichotomy between explicit and implicit representations is less fundamental than it initially appeared, and that the most effective representations will combine elements of both.

5.4 Limitations

Four limitations apply. First, evaluations used single-scene methods trained per-scene; generalisation models trained across scenes were not evaluated. Second, training hardware was fixed at a single A100; consumer GPUs show different relative performance. Third, dynamic scenes (moving objects) were not evaluated; all methods assume static scenes. Fourth, the study evaluated single-object or single-scene reconstruction; city-scale or building-scale methods require different architectures.

6. Conclusion

6.1 Summary

Evaluation of five neural rendering methods confirms: 3DGS achieves highest Blender PSNR (33.2 dB) and real-time rendering (68.4 FPS). Instant-NGP trains in 2.4 minutes -- 300x faster than NeRF. Mip-NeRF 360 achieves highest unbounded scene quality (29.4 dB). NeuS achieves 3x better surface reconstruction than all view-synthesis methods (Chamfer = 0.62 mm).

6.2 Recommendations

Three recommendations emerge. First, use 3DGS as the default for interactive applications requiring real-time rendering -- it achieves the best quality-speed trade-off. Second, use Instant-NGP for rapid capture workflows where 2-5 minute

training provides immediate feedback for quality assessment. Third, use NeuS exclusively for applications requiring accurate surface geometry -- its 3x Chamfer distance advantage over 3DGS is irreplaceable for metrology, heritage documentation, and medical imaging.

References

- Barron, J.T., Mildenhall, B., Tancik, M., Hedman, P., Martin-Brualla, R. and Srinivasan, P.P. (2021) Mip-NeRF: A multiscale representation for anti-aliasing neural radiance fields. *Proceedings of ICCV 2021*, pp. 5855-5864.
- Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P. and Hedman, P. (2022) Mip-NeRF 360: Unbounded anti-aliased neural radiance fields. *Proceedings of CVPR 2022*, pp. 5470-5479.
- Boss, M., Braun, R., Jampani, V., Barron, J.T., Liu, C. and Lensch, H. (2021) NeRD: Neural reflectance decomposition from image collections. *Proceedings of ICCV 2021*, pp. 12684-12694.
- Chen, A., Xu, Z., Geiger, A., Yu, J. and Su, H. (2022) TensorRF: Tensorial radiance fields. *Proceedings of ECCV 2022*, pp. 333-350.
- Hedman, P., Philip, J., Price, T., Frahm, J.M., Drettakis, G. and Brostow, G. (2018) Deep blending for free-viewpoint image-based rendering. *ACM Transactions on Graphics*, 37(6), pp. 1-15.
- Jensen, R., Dahl, A., Vogiatzis, G., Tola, E. and Aanaes, H. (2014) Large scale multi-view stereopsis evaluation. *Proceedings of CVPR 2014*, pp. 406-413.
- Kerbl, B., Kopanas, G., Leimkuehler, T. and Drettakis, G. (2023) 3D Gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), pp. 1-14.
- Levoy, M. and Hanrahan, P. (1996) Light field rendering. *Proceedings of SIGGRAPH 1996*, pp. 31-42.
- Martin-Brualla, R., Radwan, N., Sajjadi, M.S.M., Barron, J.T., Dosovitskiy, A. and Duckworth, D. (2021) NeRF in the wild: Neural radiance fields for unconstrained photo collections. *Proceedings of CVPR 2021*, pp. 7210-7219.
- Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R. and Ng, R. (2020) NeRF: Representing scenes as neural radiance fields for view synthesis. *Proceedings of ECCV 2020*, pp. 405-421.
- Muller, T., Evans, A., Schied, C. and Keller, A. (2022) Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics*, 41(4), pp. 1-15.
- Park, K., Sinha, U., Barron, J.T. et al. (2021) Nerfies: Deformable neural radiance fields. *Proceedings of ICCV 2021*, pp. 5865-5874.

- Pumarola, A., Corona, E., Pons-Moll, G. and Moreno-Noguer, F. (2021) D-NeRF: Neural radiance fields for dynamic scenes. *Proceedings of CVPR 2021*, pp. 10318-10327.
- Schonberger, J.L. and Frahm, J.M. (2016) Structure-from-motion revisited. *Proceedings of CVPR 2016*, pp. 4104-4113.
- Verbin, D., Hedman, P., Mildenhall, B., Zickler, T., Barron, J.T. and Srinivasan, P.P. (2022) Ref-NeRF: Structured view-dependent appearance for neural radiance fields. *Proceedings of CVPR 2022*, pp. 5481-5490.
- Wang, P., Liu, L., Liu, Y., Theobalt, C., Komura, T. and Wang, W. (2021) NeuS: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *Advances in Neural Information Processing Systems*, 34, pp. 27171-27183.
- Yariv, L., Gu, J., Kasten, Y. and Lipman, Y. (2021) Volume rendering of neural implicit surfaces. *Advances in Neural Information Processing Systems*, 34, pp. 4805-4815.
- Zhang, K., Riegler, G., Snavely, N. and Koltun, V. (2020) NeRF++: Analyzing and improving neural radiance fields. *arXiv preprint arXiv:2010.07492*.

Declarations

Funding

This research was supported by the Italian Ministry of University under grant PRIN-2022-NeuralRender. Computational resources provided by CINECA Leonardo supercomputer. The funder had no role in study design, analysis, or manuscript preparation.

Conflict of Interest

The author declares no conflicts of interest.

Data Availability Statement

All benchmark evaluation scripts and results are deposited in the Zenodo repository under DOI: 10.5281/zenodo.19614730. NeRF Synthetic, LLFF, Mip-NeRF 360, and DTU datasets are publicly available. Released under MIT License (code) and CC BY 4.0 (results).

Ethical Approval

This study involved computational experiments on publicly available 3D scene datasets. No human subjects, personal data, or proprietary scenes were involved. No institutional ethical approval was required.

Appendix A

Neural Rendering Method Selection Guide

This appendix provides a decision tree for selecting neural rendering methods based on application requirements and computational constraints.

Step 1: Primary Application

- 1a. Real-time interactive (VR, games, digital twin) --> 3D Gaussian Splatting (68.4 FPS; best quality)
- 1b. Rapid capture / e-commerce preview --> Instant-NGP (2.4 min training; immediate feedback)
- 1c. High-quality unbounded outdoor scene --> Mip-NeRF 360 (best outdoor PSNR; hours training)
- 1d. Surface mesh needed (metrology, heritage, medical) --> NeuS (0.62 mm Chamfer; only accurate option)

Step 2: Computational Budget

- 2a. < 10 minutes available --> Instant-NGP only (2.4 min to convergence)
- 2b. 30-60 minutes --> 3DGS (30-60 min; best quality-speed balance)
- 2c. Hours available, quality critical --> Mip-NeRF 360 (unbounded) or NeuS (surface accuracy)

Step 3: Scene Type

- 3a. Bounded object or room --> Any method; 3DGS or Instant-NGP recommended
- 3b. Forward-facing (single direction) --> Instant-NGP or Mip-NeRF 360
- 3c. Unbounded 360-degree outdoor --> Mip-NeRF 360 (contract mapping handles infinite space)
- 3d. Dynamic scene (moving objects) --> None of these; use dynamic NeRF variants