NEURAL INFORMATION PROCESSING SYSTEMS

GSDF: 3DGS Meets SDF for Improved Neural Rendering and Reconstruction

Mulin Yu, Tao Lu, Linning Xu, Lihan Jiang, Yuanbo Xiangli, Bo Dai

上海人工智能实验室
Shanghai Artificial Intelligence Laboratory

Problem Statement

Input: Multi-view images and their poses

Problem Statement

Input: Multi-view images and their poses

Output: A reconstructed model with accurate **geometry** and realistic **rendering** quality

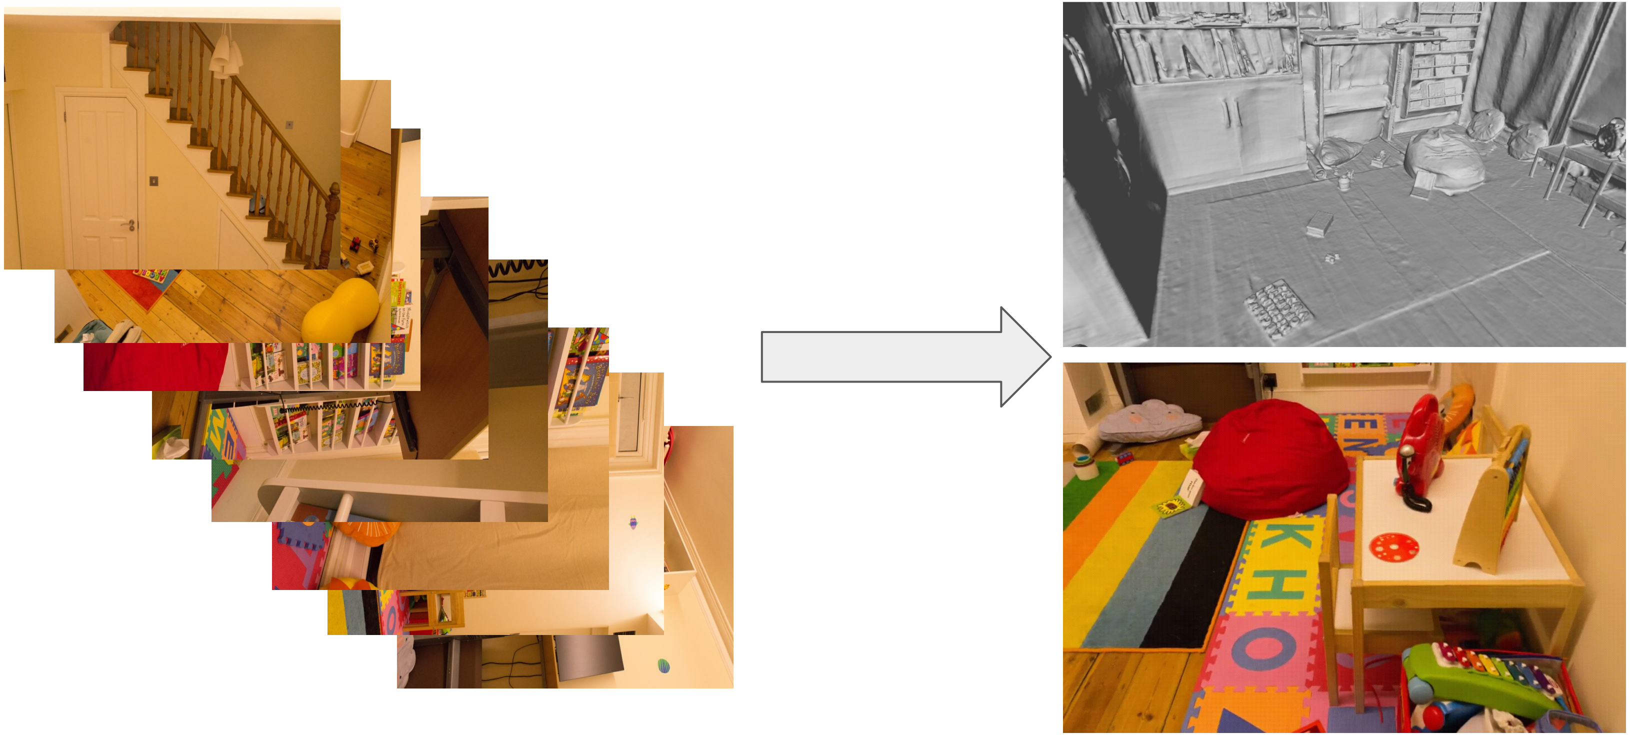

Related Work

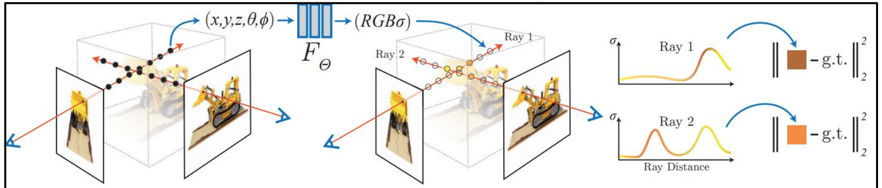

Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In ECCV, 2020

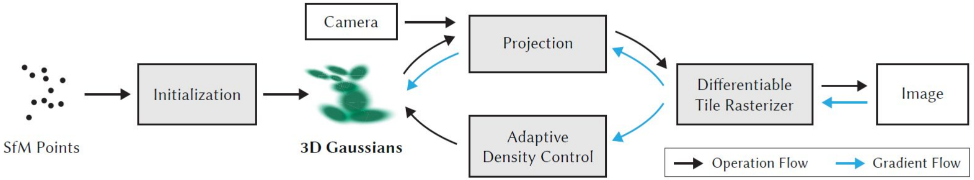

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field endering. ACM Transactions on Graphics, 42(4), 2023.

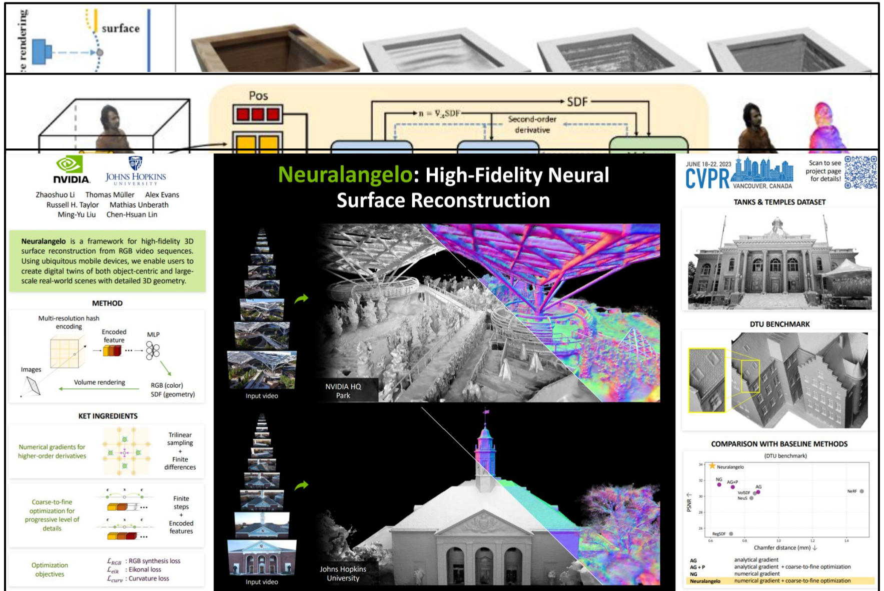

Zhaoshuo Li, Thomas Müller, Alex Evans, Russell H Taylor, Mathias Unberath, Ming-Yu Liu, and Chen-Hsuan Lin. Neuralangelo: High-fidelity neural surface reconstruction. CVPR, pages 8456–8465,2023.

Antoine Guédon and Vincent Lepetit. Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering. arXiv preprint arXiv:2311.12775, 2023.

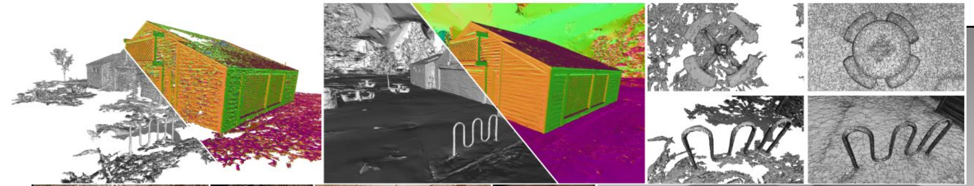

Ze ... Gaussian opacity fi ... nd compact surface reconstruction in unbounded scenes. arXiv preprint arXiv:2404.10772 (2024)

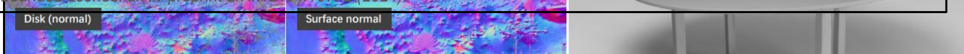

Related Work

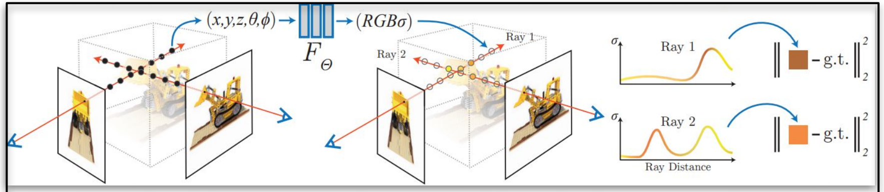

Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In ECCV, 2020

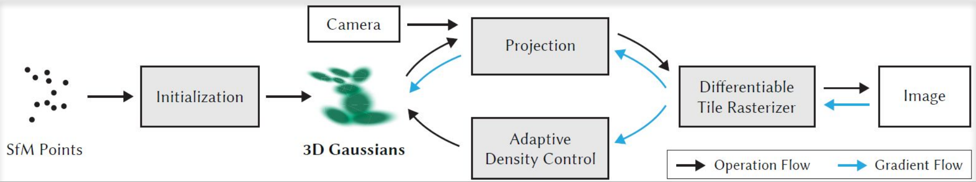

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field endering. ACM Transactions on Graphics, 42(4), 2023.

Related Work

Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In ECCV, 2020

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field endering. ACM Transactions on Graphics, 42(4), 2023.

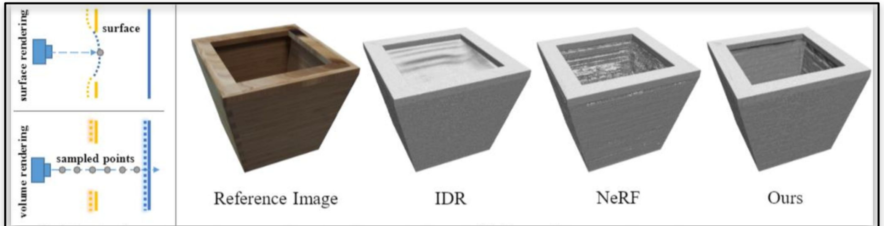

Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. arXiv preprint arXiv:2106.10689, 2021.

Related Work

Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In ECCV, 2020

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field endering. ACM Transactions on Graphics, 42(4), 2023.

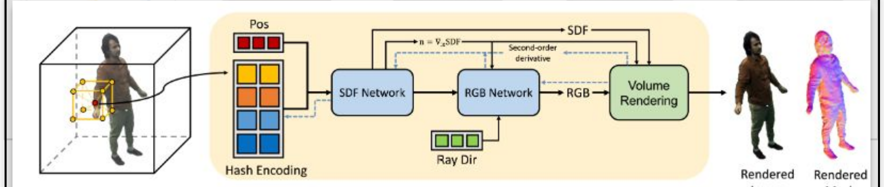

(a) NeuS2: static scene reconstruction

Yiming Wang, Qin Han, Marc Habermann, Kostas Daniilidis, Christian Theobalt, and Lingjie Liu. Neus2: Fast learning of neural implicit surfaces for multi-view reconstruction. 2023 ICCV, pages 3272–3283, 2022.

Related Work

Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In ECCV, 2020

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field endering. ACM Transactions on Graphics, 42(4), 2023.

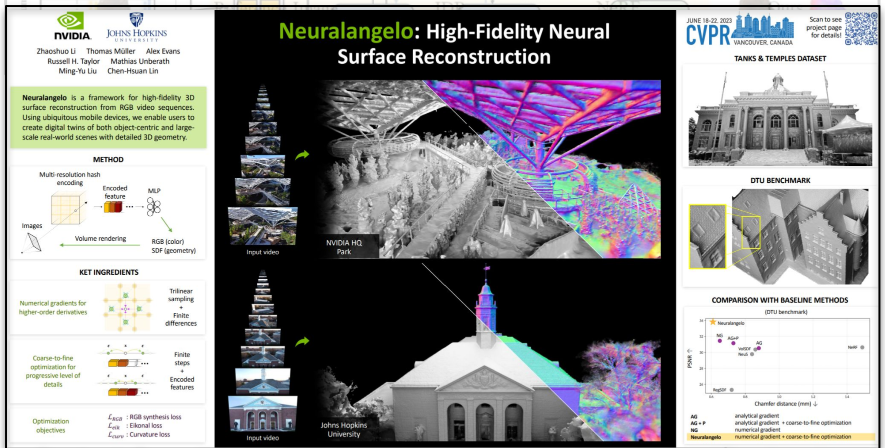

Zhaoshuo Li, Thomas Müller, Alex Evans, Russell H Taylor, Mathias Unberath, Ming-Yu Liu, and Chen-Hsuan Lin. Neuralangelo: High-fidelity neural surface reconstruction. CVPR, pages 8456–8465,2023.

Related Work

Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In ECCV, 2020

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field endering. ACM Transactions on Graphics, 42(4), 2023.

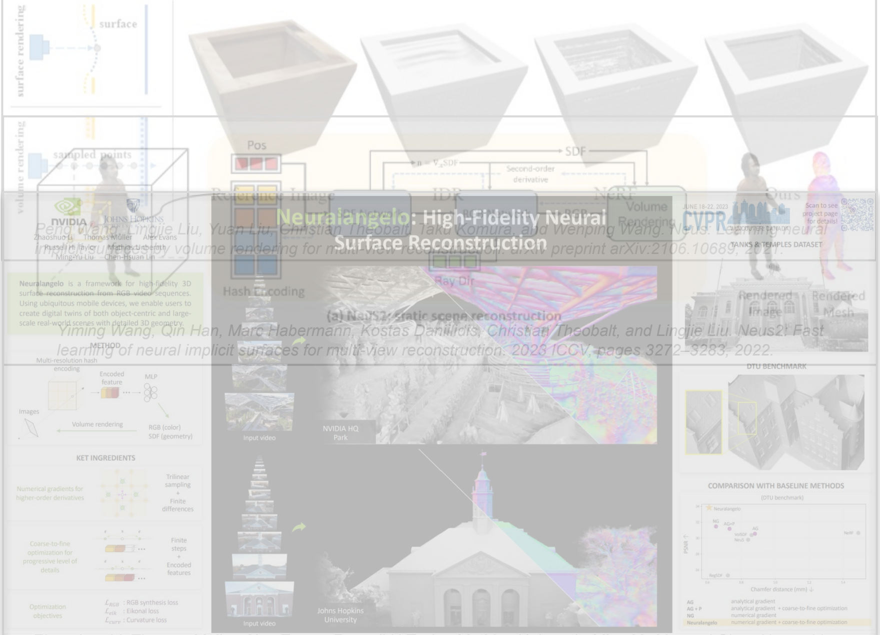

Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. arXiv preprint arXiv:2106.10689, 2021.

Yiming Wang, Qin Han, Marc Habermann, Kostas Daniilidis, Christian Theobalt, and Lingjie Liu. Neus2: Fast learning of neural implicit surfaces for multi-view reconstruction. 2023 ICCV, pages 3272–3283, 2022.

Zhaoshuo Li, Thomas Müller, Alex Evans, Russell H Taylor, Mathias Unberath, Ming-Yu Liu, and Chen-Hsuan Lin. Neuralangelo: High-fidelity neural surface reconstruction. CVPR, pages 8456–8465,2023.

Antoine Guédon and Vincent Lepetit. Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering. arXiv preprint arXiv:2311.12775, 2023.

Related Work

Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In ECCV, 2020

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field endering. ACM Transactions on Graphics, 42(4), 2023.

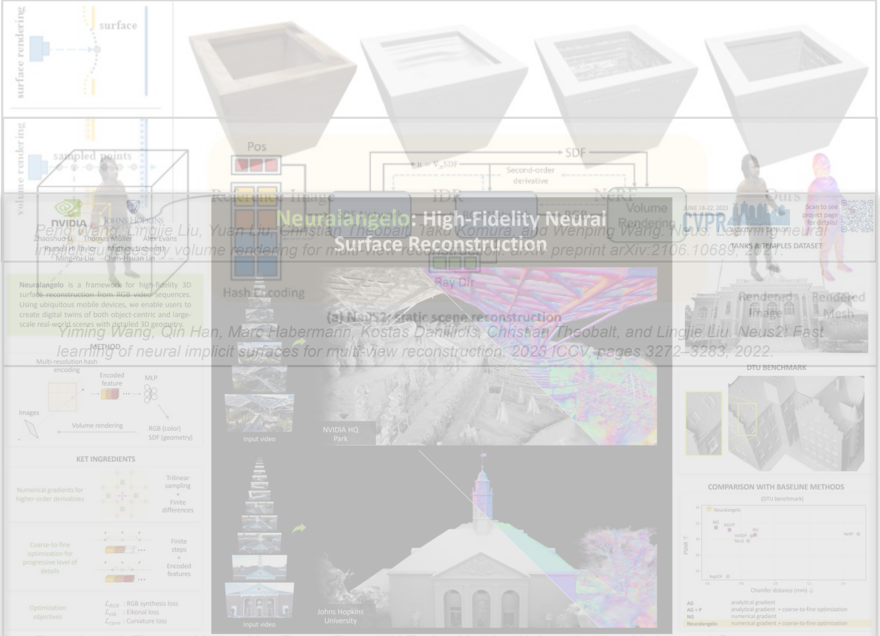

Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. arXiv preprint arXiv:2106.10689, 2021.

Yiming Wang, Qin Han, Marc Habermann, Kostas Daniilidis, Christian Theobalt, and Lingjie Liu. Neus2: Fast learning of neural implicit surfaces for multi-view reconstruction. 2023 ICCV, pages 3272–3283, 2022.

Zhaoshuo Li, Thomas Müller, Alex Evans, Russell H Taylor, Mathias Unberath, Ming-Yu Liu, and Chen-Hsuan Lin. Neuralangelo: High-fidelity neural surface reconstruction. CVPR, pages 8456–8465,2023.

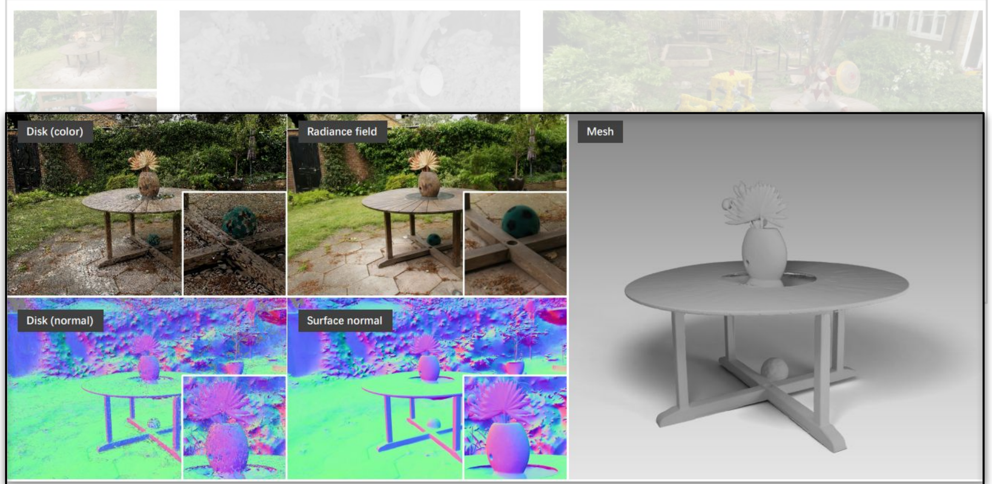

Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2D gaussian splatting for geometrically accurate radiance fields. ArXiv, abs/2403.17888, 2024.

Related Work

Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In ECCV, 2020

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field endering. ACM Transactions on Graphics, 42(4), 2023.

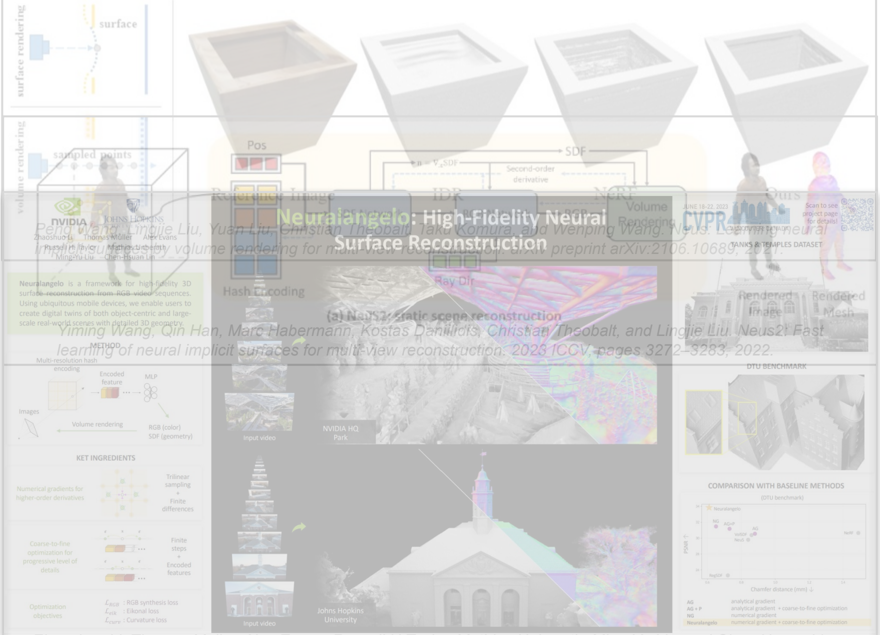

Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. arXiv preprint arXiv:2106.10689, 2021.

Yiming Wang, Qin Han, Marc Habermann, Kostas Daniilidis, Christian Theobalt, and Lingjie Liu. Neus2: Fast learning of neural implicit surfaces for multi-view reconstruction. 2023 ICCV, pages 3272–3283, 2022.

Zhaoshuo Li, Thomas Müller, Alex Evans, Russell H Taylor, Mathias Unberath, Ming-Yu Liu, and Chen-Hsuan Lin. Neuralangelo: High-fidelity neural surface reconstruction. CVPR, pages 8456–8465,2023.

Antoine Guédon and Vincent Lepetit. Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering. arXiv preprint arXiv:2311.12775, 2023.

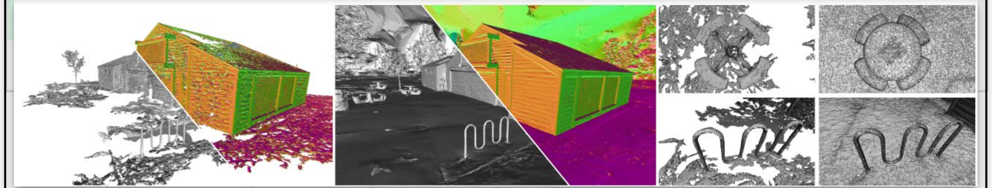

Zehao Yu, Sattler Torsten, and Geiger Andreas. Gaussian opacity fields: Efficient and compact surface reconstruction in unbounded scenes. arXiv preprint arXiv:2404.10772 (2024).

Related Work

Current methods impose strict constraints on density fields or primitive shapes, which enhances the affinity for geometric reconstruction at the sacrifice of rendering quality.

Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In ECCV, 2020

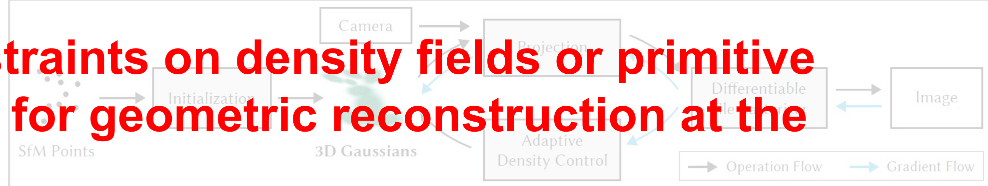

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field endering. ACM Transactions on Graphics, 42(4), 2023.

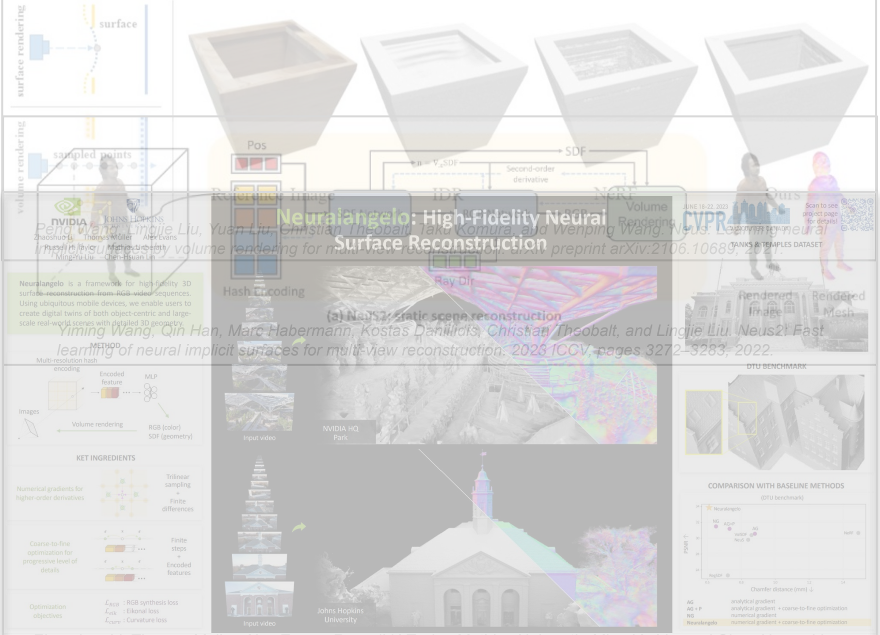

Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. arXiv preprint arXiv:2106.10689, 2021.

Yiming Wang, Qin Han, Marc Habermann, Kostas Daniilidis, Christian Theobalt, and Lingjie Liu. Neus2: Fast learning of neural implicit surfaces for multi-view reconstruction. 2023 ICCV, pages 3272–3283, 2022.

Zhaoshuo Li, Thomas Müller, Alex Evans, Russell H Taylor, Mathias Unberath, Ming-Yu Liu, and Chen-Hsuan Lin. Neuralangelo: High-fidelity neural surface reconstruction. CVPR, pages 8456–8465,2023.

Antoine Guédon and Vincent Lepetit. Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering. arXiv preprint arXiv:2311.12775, 2023.

Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2D gaussian splatting for geometrically accurate radiance fields. ArXiv, abs/2403.17888, 2024.

Zehao Yu, Sattler Torsten, and Geiger Andreas. Gaussian opacity fields: Efficient and compact surface reconstruction in unbounded scenes. arXiv preprint arXiv:2404.10772 (2024).

Two-branch Framework

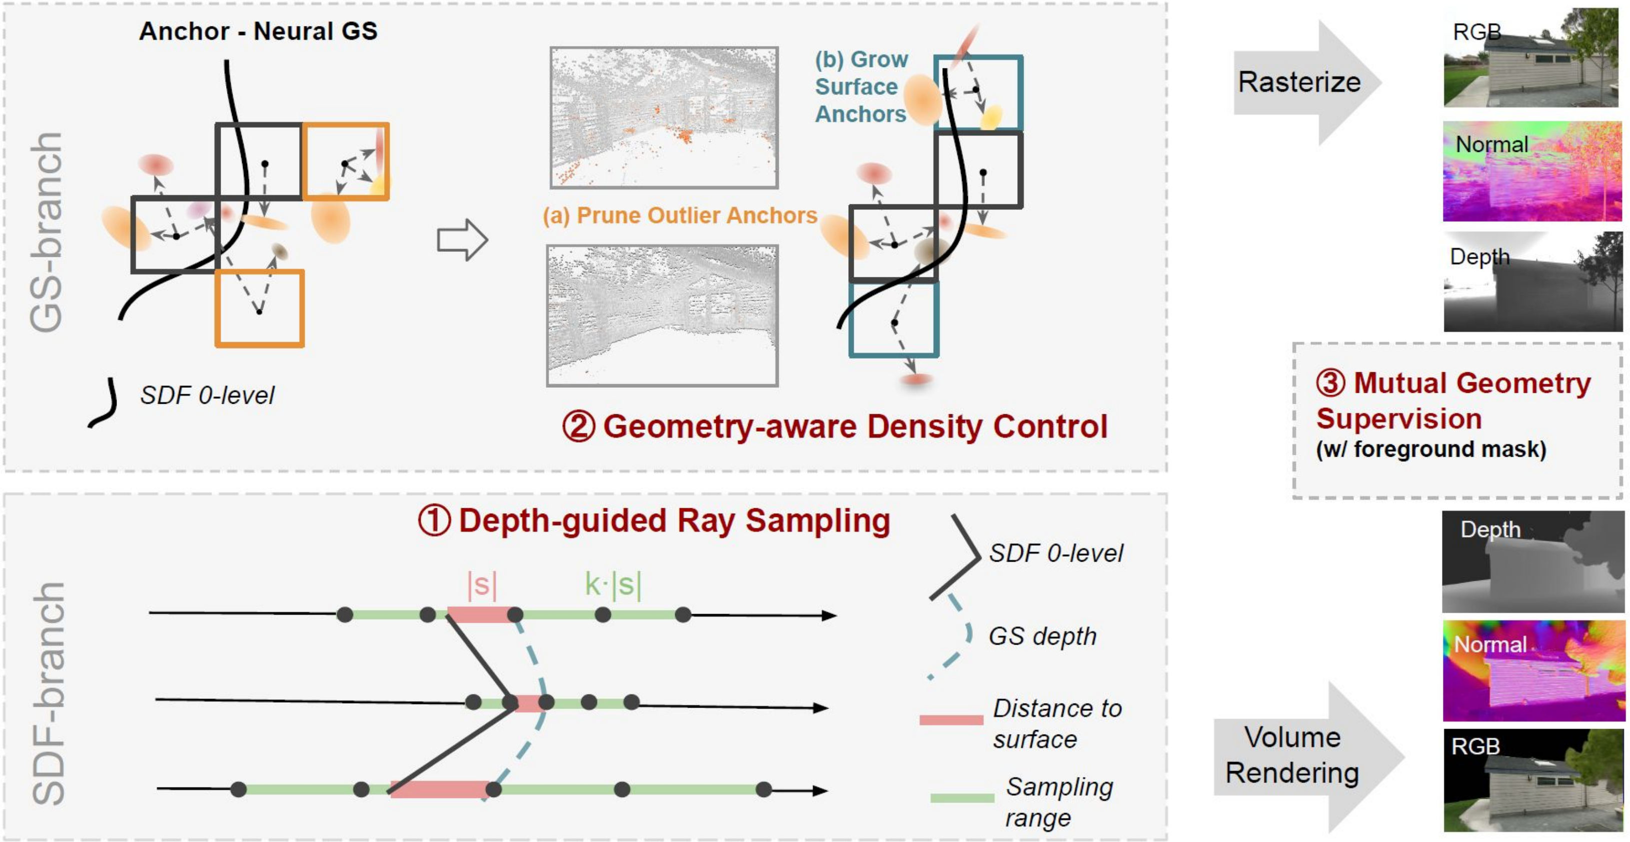

GS → SDF: Depth Guided Ray Sampling

We leverage depth maps from the GS-branch to refine the ray sampling range for the SDF-branch.

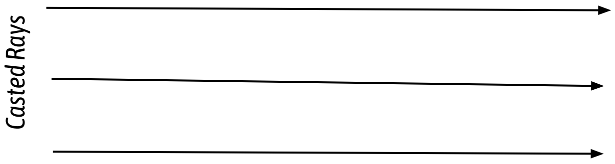

GS → SDF: Depth Guided Ray Sampling

We leverage depth maps from the GS-branch to refine the ray sampling range for the SDF-branch.

$$D = \sum_{i \in N} d_i \sigma_i \prod_{j=1}^{i-1} (1 - \sigma_j)$$

GS depth

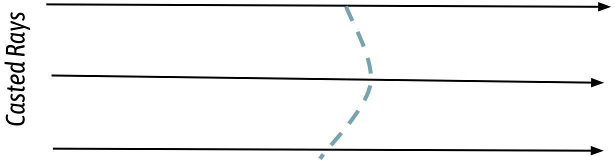

GS depth

GS → SDF: Depth Guided Ray Sampling

We leverage depth maps from the GS-branch to refine the ray sampling range for the SDF-branch.

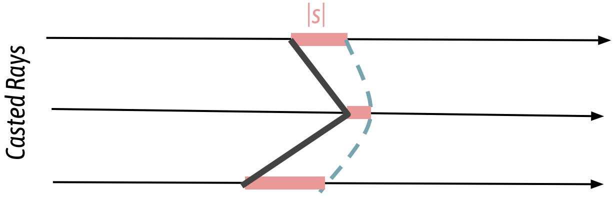

$$D = \sum_{i \in N} d_i \sigma_i \prod_{j=1}^{i-1} (1 - \sigma_j)$$

GS depth

$$s = \mathcal{F}_s(\vec{o} + D \cdot \vec{v})$$

SDF 0-level

GS depth

GS → SDF: Depth Guided Ray Sampling

We leverage depth maps from the GS-branch to refine the ray sampling range for the SDF-branch.

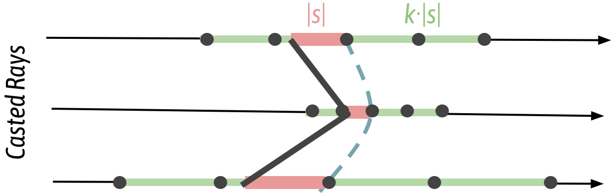

$$D = \sum_{i \in N} d_i \sigma_i \prod_{j=1}^{i-1} (1 - \sigma_j)$$

GS depth

$$s = \mathcal{F}_s(\vec{o} + D \cdot \vec{v})$$

SDF 0-level

GS depth

SDF → GS: Geometry-aware Gaussian Density Control

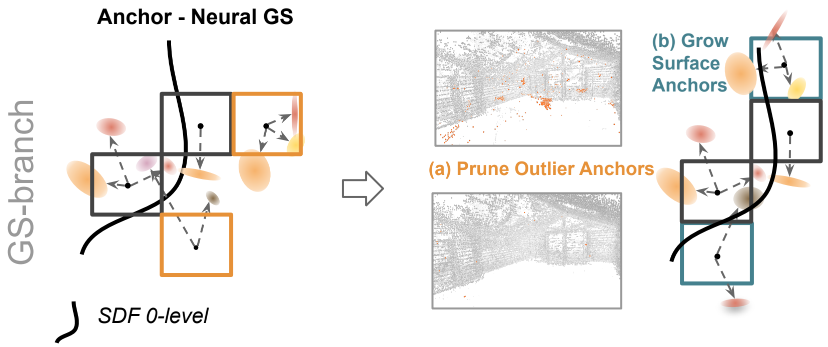

SDF $\rightarrow$ GS: Geometry-aware Gaussian Density Control

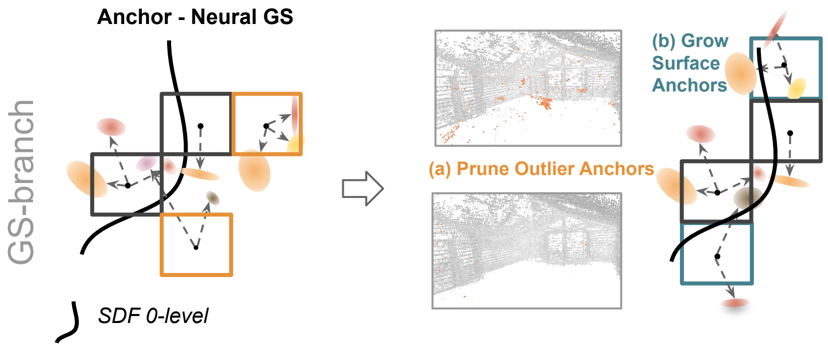

$$\mu(s) = \exp\left(-s^2/(2\sigma^2)\right)$$

SDF → GS: Geometry-aware Gaussian Density Control

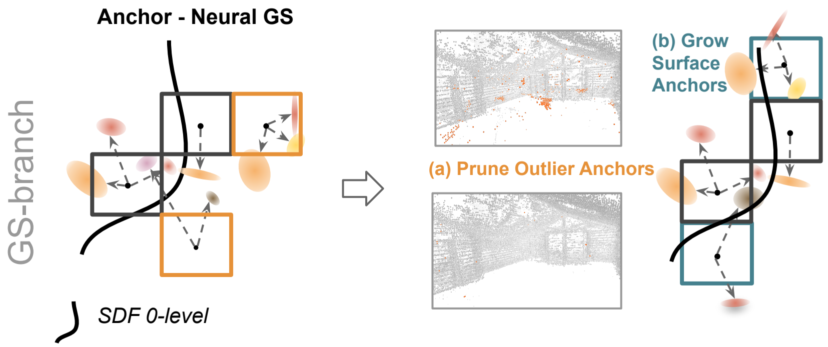

$$\mu(s) = \exp\left(-s^2/(2\sigma^2)\right)$$

$$\epsilon_g = \nabla_g + \omega_g \mu(s)$$

$$\epsilon_p = \sigma_a - \omega_p(1 - \mu(s))$$

GS $\leftrightarrow$ SDF: Mutual Geometry Supervision

$$\mathcal{L}_{\text{mutual}} = \lambda_d \mathcal{L}_d + \lambda_n \mathcal{L}_n = \lambda_d \|D_{gs} - D_s\| + \lambda_n \left(1 - \frac{|N_{gs} \cdot N_s|}{\|N_{gs}\| \, \|N_s\|}\right)$$

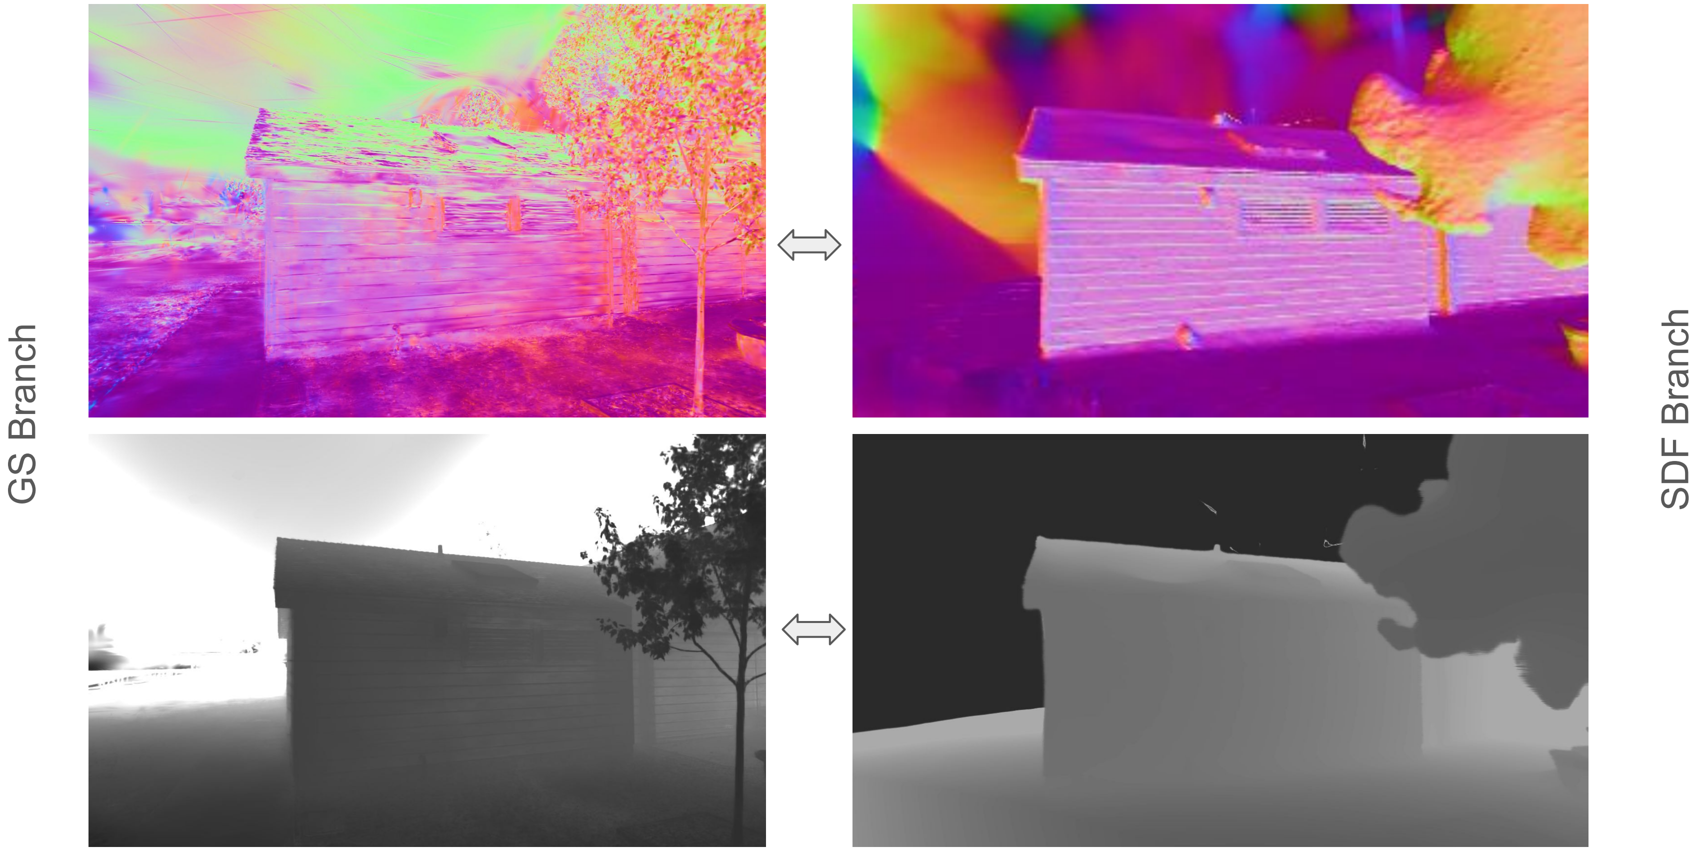

Comparison

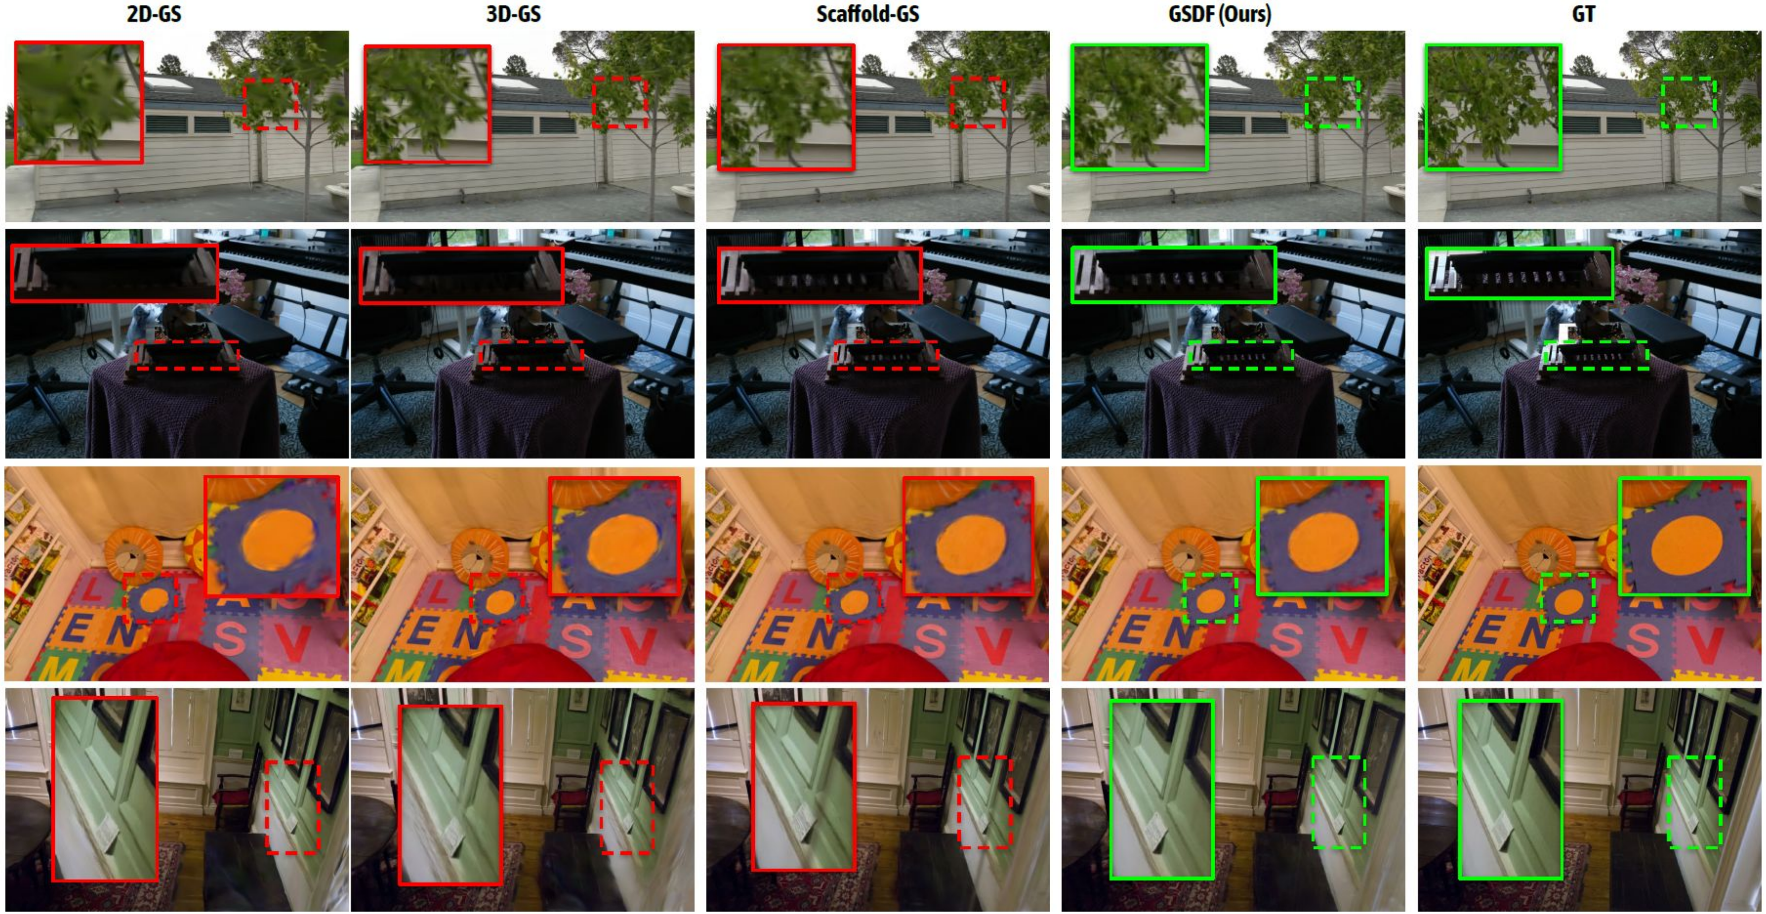

Comparison - Random initialization

Comparison

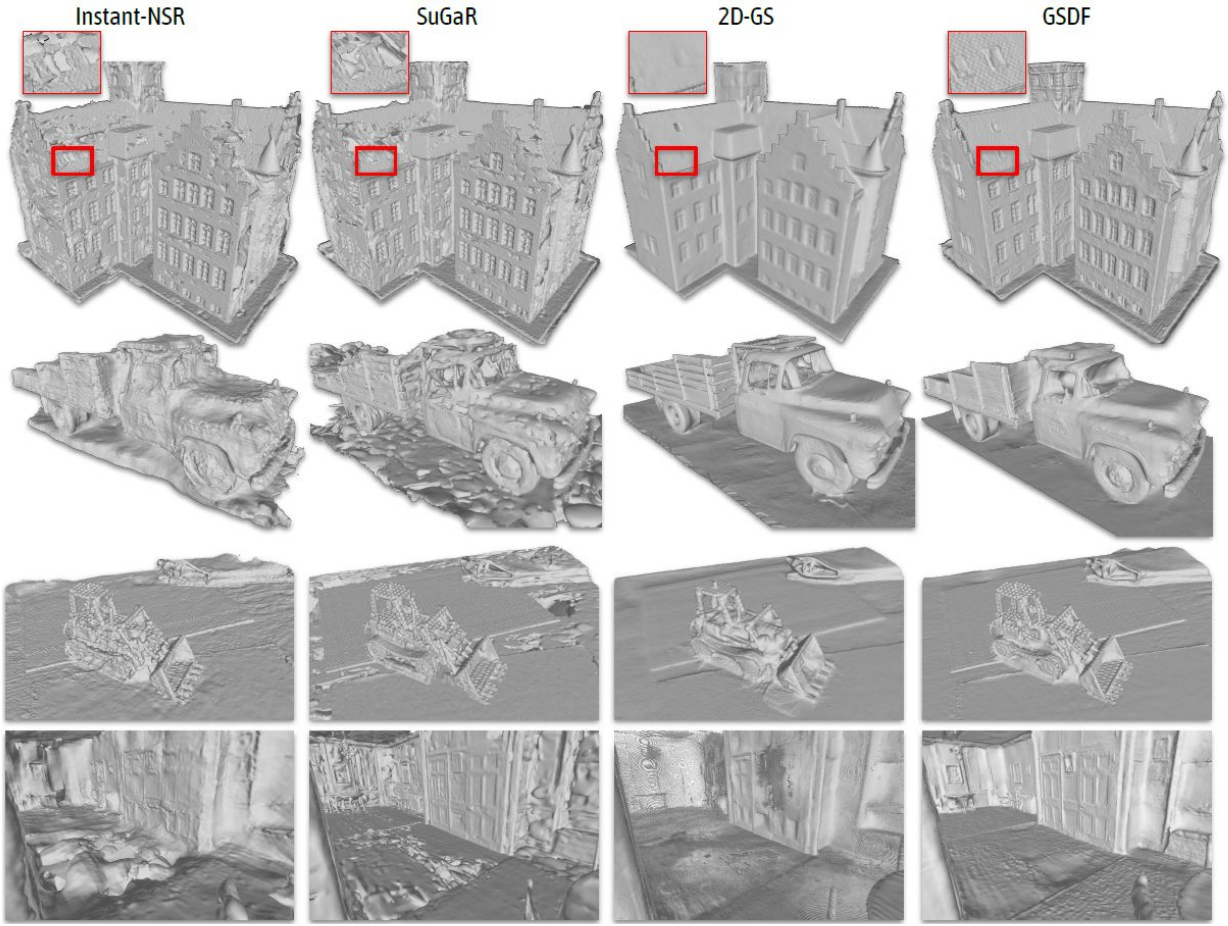

Comparison

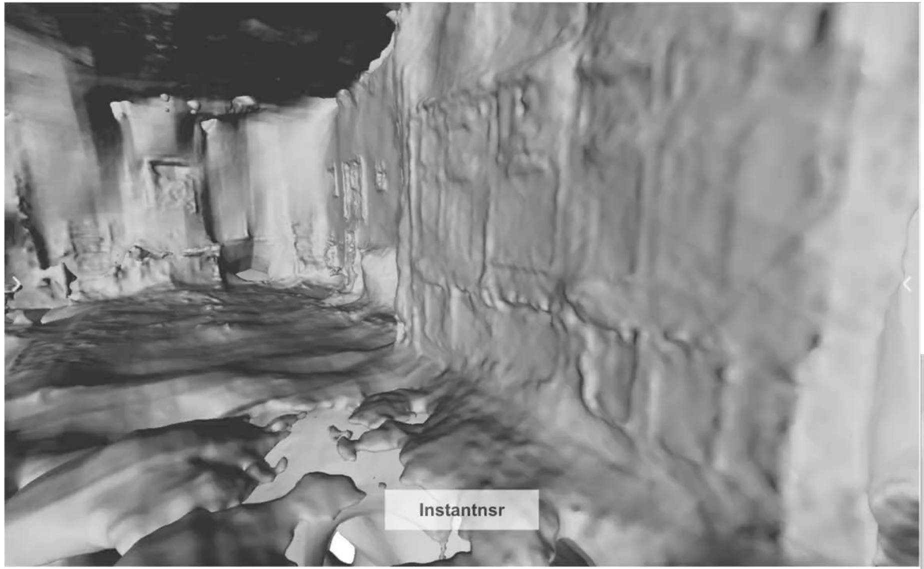

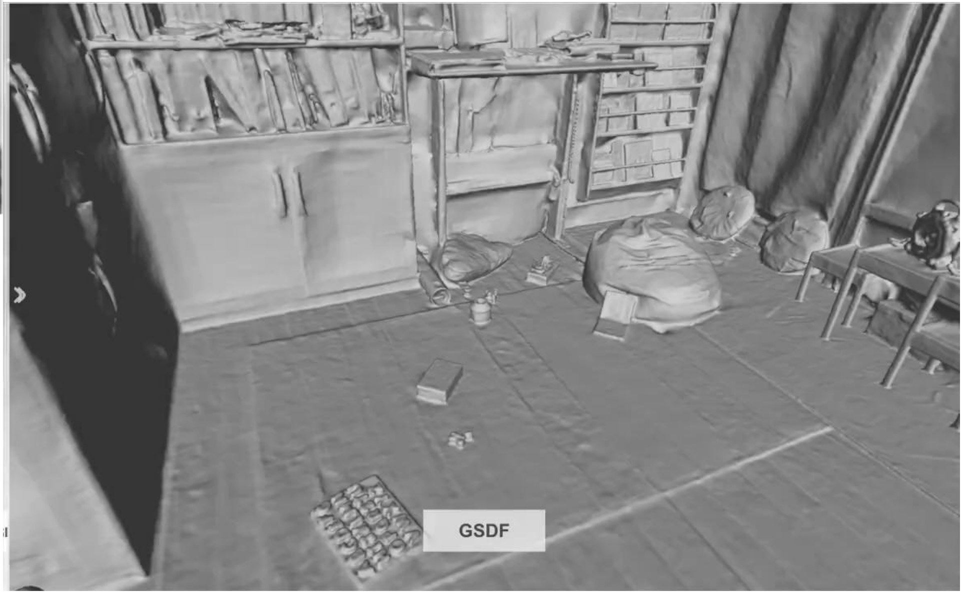

Thanks!

https://city-super.github.io/GSDF/