

Groupe de lecture :

SPLAT THE NET :

RADIANCE FIELDS WITH SPLATTABLE

NEURAL PRIMITIVES



Buonomo Camille,
17/12/2025,

Xilong Zhou*, Bao-Huy Nguyen*, Loic Magne, Vladislav Golyanik, Thomas Leimkuhler (slides), Christian Theobal ¹

Virtual Scenes are Essential for...




[wilo.com]

Computer Graphics
e.g., Virtual Reality




[ce.cit.tum.de]

Computer Vision & AI
e.g., Robotics



[developer.nvidia.com]

Medicine
e.g., Surgical Planning



[Xie et al., 2025]

Earth Sciences
e.g., Climate Simulations



[pompeiiisites.org]


Archaeology
e.g., Masonry Analysis

...


Reconstruction from Observations




Virtual Scenes and Images



Virtual Scenes and Images




Scene Representation Zoo



Mesh



Voxel Grid




Neural Field

⚠ Topology difficult to optimize


⚠ Making rendering differentiable is hard

⚠ Too much memory


Neural Fields




Image




Video



Signed Distance Field




Sound





Radiance Field

...

Neural Fields




Neural Radiance Fields




Mildenhall, Srinivasan, Tancik, Barron, Ramamoorthi, Ng. **NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis**. ECCV 2020.

Virtual Scenes and Images



Emission—Absorption Rendering




Pixel Color

$$C = \int_{t_{\text{near}}}^{t_{\text{far}}} \underbrace{T(t)}_{\text{Transmittance}} \underbrace{\sigma(\mathbf{x}(t))}_{\text{Vol. Density}} \underbrace{c(\mathbf{x}(t), \mathbf{d})}_{\text{Color}} dt$$

$$T(t) = \exp \left(- \int_{t_{\text{near}}}^t \underbrace{\sigma(\mathbf{x}(s))}_{\text{Vol. Density}} ds \right)$$

Kajiya, Von Herzen. **Ray Tracing Volume Densities**. SIGGRAPH 1984.

Emission—Absorption Rendering



$$C = \sum_i \overset{\text{Transmittance}}{T_i} \overset{\text{Vol. Density}}{(1 - \exp(-\sigma_i \delta_i))} \overset{\text{Color}}{c_i}$$

Pixel Color

$$T_i = \exp \left(- \sum_{j < i} \sigma_j \delta_j \right)$$

Max. Optical Models for Direct Volume Rendering. TVCG 1995.

Rendering Algorithm

For each pixel:

Sample locations on view ray

Evaluate network for each sample

Perform volume rendering:
$$C = \sum_i \exp\left(-\sum_{j<i} \sigma_j \delta_j\right) (1 - \exp(-\sigma_i \delta_i)) c_i$$


Training

5D Input
Position + Direction

4D Output
Density + Color

Volume
Rendering

Rendering
Loss




[Mildenhall et al., 2020]

Improvements & Extensions


- Architectures

- Positional encoding [Mildenhall et al. 2020, Tancik et al. 2020, ...]
- Explicit and hybrid models [Mueller et al. 2022, Fridovich-Keil et al. 2022, ...]
- ...



- Training schemes


- Coarse-to-fine and multiscale [Mildenhall et al. 2020, Barron et al. 2022, ...]
- Efficient ray sampling [Deng et al. 2022, Lombardi et al. 2021, ...]
- ...




- Higher dimensions

- Dynamic scenes [Park et al. 2021, Tretschk et al. 2021, ...]
- Relighting [Srinivasan et al. 2021, Martin-Brualla et al. 2021, ...]
- ...


- ...




Neural Radiance Fields are Compact




5 MB




Neural Radiance Fields are Inefficient



Speed vs. Quality




Mueller, Evans, Schied, Keller. **Instant Neural Graphics Primitives with a Multiresolution Hash Encoding**. SIGGRAPH 2022.



Barron, Mildenhall, Verbin, Srinivasan, Hedman. **Mip-NeRF 360: Unbounded Anti-Aliased Neural Radiance Fields**. CVPR 2022.


Scene Representation Zoo




Mesh



Voxel Grid



Neural Field



Unstructured Primitives


⚠ Topology difficult to optimize

⚠ Making rendering differentiable is hard


⚠ Too much memory

⚠ Too slow


Point-based Graphics



[Levoy & Whitted 1985]




[Max 2004]




[Bergmeir et al. 2016]

Virtual Scenes and Images





Gaussian Mixture Models



[github.com/DandiMahendris]

Scene as 3D Gaussian Mixture Model



Volumetric
Density


3D Gaussian

$$\sigma(\mathbf{x}) = \sum_i G^{(3D)}(\mathbf{x}; \underbrace{\boldsymbol{\mu}_i}_{\text{Mean}}, \underbrace{\boldsymbol{\Sigma}_i}_{\text{Covariance}}, \underbrace{o_i}_{\text{Opacity}})$$


Additionally:

Each Gaussian has a color $\mathbf{c}$,
which is a function of direction $\mathbf{d}$

Virtual Scenes and Images



Emission—Absorption Rendering




Pixel Color

$$C = \int_{t_{\text{near}}}^{t_{\text{far}}} \underbrace{T(t)}_{\text{Transmittance}} \underbrace{\sigma(\mathbf{x}(t))}_{\text{Vol. Density}} \underbrace{c(\mathbf{x}(t), \mathbf{d})}_{\text{Color}} dt$$

$$T(t) = \exp \left(- \int_{t_{\text{near}}}^t \underbrace{\sigma(\mathbf{x}(s))}_{\text{Vol. Density}} ds \right)$$

Kajiya, Von Herzen. **Ray Tracing Volume Densities**. SIGGRAPH 1984.

Splatting




$$G^{(3D)} \xrightarrow{\text{projection}} G^{(2D)}$$

(under mild assumptions –
I am skipping a couple of details here)

Zwicker, Pfister, van Baar, Gross. **EWA Volume Splatting**. Visualization 2001.

Splatting



Pixel Color

Projected Gaussian Shape

Gaussian Color

$$C = \sum_i G_i^{(2D)} c_i \underbrace{\prod_{j < i} (1 - G_j^{(2D)})}_{\text{Transmittance (ensure that Gaussians occlude the ones behind)}}$$

Zwicker, Pfister, van Baar, Gross. **EWA Volume Splatting**. Visualization 2001.

Rendering Algorithm

Project Gaussians to the screen

Sort Gaussians based on depth

For each pixel:


Identify Gaussians overlapping the pixel

Splat Gaussians:
$$C = \sum_i G_i^{(2D)} c_i \prod_{j < i} (1 - G_j^{(2D)})$$


First Result




Population Control



Speed & Quality



Mueller, Evans, Schied, Keller. **Instant Neural Graphics Primitives with a Multiresolution Hash Encoding**. SIGGRAPH 2022.



3D Gaussian Splatting







Improvements & Extensions

• Model

- Anti-aliasing [Yu et al. 2023, ...]
- Structural constraints [Lu et al. 2024, ...]
- Temporal coherence [Radl et al. 2024, ...]
- Primitive types [Liu et al. 2025, ...]
- Compression [Niedermayr et al. 2024]
- Hierarchical [Kerbl et al., 2024]
- ...

• Higher dimensions


- Dynamic scenes [Luiten et al. 2024, ...]
- Relighting [Bi et al. 2024, ...]
- ...




3D Gaussians are Not Compact



A Gaussian has
limited expressive power




Scenes routinely require
millions of Gaussians



Typically:
> 1 GB per scene


Scene Representation Zoo




Mesh



Voxel Grid



Neural Field



Unstruct. Primitives

⚠ Topology difficult to optimize

⚠ Making rendering differentiable is hard

⚠ Too much memory

⚠ Too slow

⚠ Limited expressivity




Scene Representation Zoo

Quality





Speed

Expressivity

Neural Field



Unstruct. Primitives




Scene Representation Zoo

Quality


Speed

Expressivity


Neural Field



Unstruct. Primitives




Splatting of
Unstructured Primitives



Neural
Representation


Neural Primitives




Gaussian




Neural



How to Splat a Neural Primitive?





Integrating a Shallow Network



$$f_{\mathbf{w}}(\begin{smallmatrix} \square \\ \square \\ \square \\ \square \\ \square \\ \square \\ \square \\ \square \end{smallmatrix} \mathbf{x}) = \begin{smallmatrix} \square & \square & \square & \square & \square & \square & \square & \square & \square & \square \end{smallmatrix} \mathbf{w}_2 \cos \left(\begin{smallmatrix} \square & \square & \square & \square \\ \square & \square & \square & \square \\ \square & \square & \square & \square \\ \square & \square & \square & \square \\ \square & \square & \square & \square \end{smallmatrix} \begin{smallmatrix} \square \\ \square \\ \square \\ \square \end{smallmatrix} + \begin{smallmatrix} \square \\ \square \\ \square \\ \square \\ \square \\ \square \\ \square \\ \square \end{smallmatrix} \right) + \begin{smallmatrix} \square \\ \square \\ \square \\ \square \end{smallmatrix} b_2$$

$\mathbf{W}_1$ $\mathbf{b}_1$


1 2 3 4



Lloyd, Irani, Ahmadi. **Using Neural Networks for Fast Numerical Integration and Optimization.** IEEE Access 2020.

Subr. **Q-NET: A Network for Low-dimensional Integrals of Neural Proxies.** Computer Graphics Forum 2021.

Integrating a Neural Primitive




$$\alpha^{(2D)} = \int_{t_{in}}^{t_{out}} \sigma(\mathbf{x}(t))$$

$$= S(t_{out}; \mathbf{r}) - S(t_{in}; \mathbf{r})$$

$$S(t; \mathbf{r}) = [W_2 \oslash W_1 \mathbf{d}]$$


$$\sin(tW_1 \mathbf{d} + W_1 \mathbf{o} + \mathbf{b}_1) + t \mathbf{b}_2$$

Gaussian Splatting



$$C = \sum_i G_i^{(2D)} \mathbf{c}_i \prod_{j < i} (1 - G_j^{(2D)})$$

Neural Primitive Splatting



$$C = \sum_i \alpha_i^{(2D)} \cancel{c_i^{(2D)}} \mathbf{c}_i \prod_{j < i} \left(1 - \alpha_j^{(2D)} \cancel{c_j^{(2D)}} \right)$$

Results



10k Primitives



Speed & Quality





3D Gaussian Splatting




Neural Splatting


Neural Primitives are Compact




Representations for Virtual Scenes




Mesh




Voxel Grid



Neural Field



Unstructured Primitives



Neural Radiance
Fields

3D Gaussian
Splatting

Neural Primitive
Splatting

The Team



Bao Hyu Nguyen



George Kopanas



Loïc Magne



Xilong Zhou

Thomas Leimkuhler



Bernhard Kerbl



Vladislav Golyanik



Christian Theobalt



George Drettakis

Thank you!



3D Gaussian Splatting for Real-time Radiance Field Rendering

Kerbl, Kopanas, Leimkuehler, Drettakis
2023



Splat the Net: Radiance Fields with Splattable Neural Primitives

Zhou, Nguyen, Magne, Golyanik, Leimkuehler, Theobalt
2025