

Physically-Based Volume Rendering and High-quality Neural Reconstruction with Volumetric Primitives

Jorge Condor
PhD Student @ USI Lugano, Switzerland


21st April 2026
Saarbrücken, Germany



(2020) BE Electronics and Control Engineer @ Unizar, Zaragoza, Spain
(2022) Masters in Graphics, Robotics and CV @ Unizar, Zaragoza, Spain
2022 – PhD Student @ USI Lugano, Switzerland
2023-2024 RS intern @ Meta Reality Labs, Zürich, Switzerland
2025 – RS intern @ NVIDIA Research, Zürich, Switzerland



Volume Rendering in Vision and Graphics



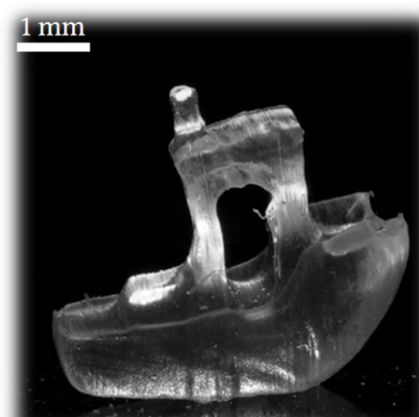
**MEDICAL AND SCIENTIFIC
VISUALIZATION**

*GSCache: Real-Time Radiance Caching for Volume Path Tracing
using 3D Gaussian Splatting*
David Bauer , Qi Wu , Hamid Gadirov , and Kwan-Liu Ma



MEDIA PRODUCTION

*Fast Volume Rendering with Spatiotemporal Reservoir
Resampling*
Daqi Lin, Chris Wyman, Cem Yuksel



**PREDICTIVE RENDERING
AND INVERSE OPTIMIZATION**

Inverse Rendering for Tomographic Volumetric Additive Manufacturing
Baptiste Nicolet, Felix Wechsler, Jorge Madrid-Wolff, Christophe Moser,
Wenzel Jakob




NOVEL-VIEW SYNTHESIS

*NeRF: Representing Scenes as Neural Radiance Fields for
View Synthesis*
Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik,
Jonathan T. Barron, Ravi Ramamoorthi, Ren Ng

Limitations of Volume Rendering


Explicit Density Fields

- Regularly spaced voxel grids, w or wo a hierarchy



Implicit Density Fields

- A neural network (MLP)




Instant Neural Graphics Primitives with a Multiresolution Hash Encoding
Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller.

Limitations of Volume Rendering

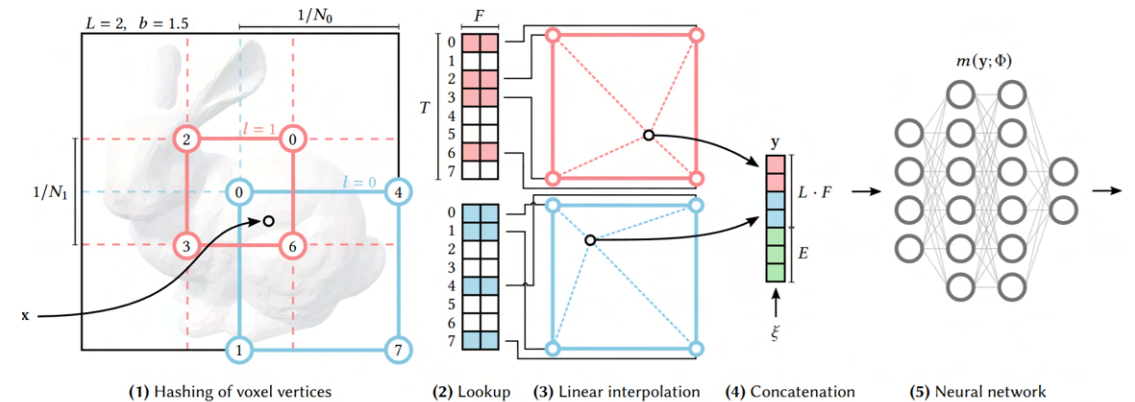
Explicit Density Fields

- Regularly spaced voxel grids, w or wo a hierarchy
 - Simple to generate (discretization) and sample from



Implicit Density Fields

- A neural network (MLP)
 - Difficult to generate (gradient-based optimization) and sample from (usually requires a separate model)




Instant Neural Graphics Primitives with a Multiresolution Hash Encoding
Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller.

Limitations of Volume Rendering

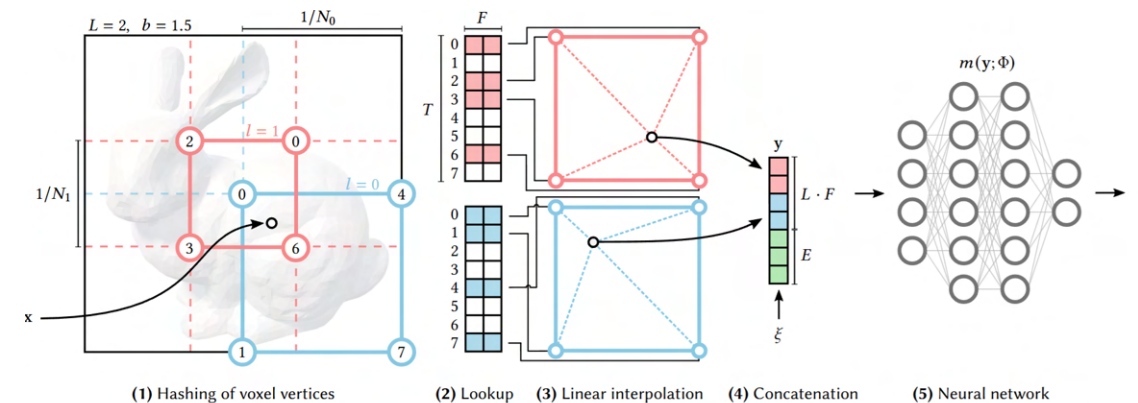
Explicit Density Fields

- Regularly spaced voxel grids, w or wo a hierarchy
 - Simple to generate (discretization) and sample from
 - Can be heavy and slow (high memory traffic)



Implicit Density Fields

- A neural network (MLP)
 - Difficult to generate (gradient-based optimization) and sample from (usually requires a separate model)
 - Very compact (unless using heavy positional encoding techniques)




Instant Neural Graphics Primitives with a Multiresolution Hash Encoding
Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller.

Limitations of Volume Rendering

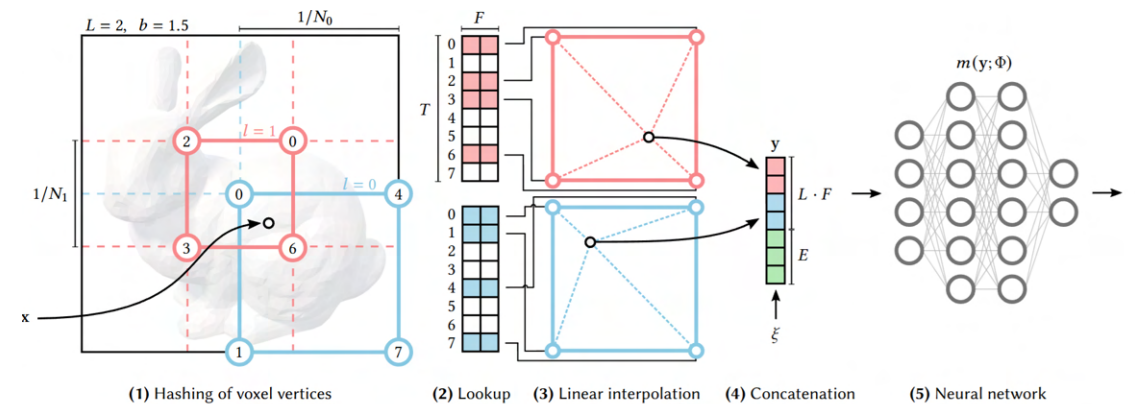
Explicit Density Fields

- Regularly spaced voxel grids, w or wo a hierarchy
 - Simple to generate (discretization) and sample from
 - Can be heavy and slow (high memory traffic)
 - Normally requires numerical methods to integrate



Implicit Density Fields

- A neural network (MLP)
 - Difficult to generate (gradient-based optimization) and sample from (usually requires a separate model)
 - Very compact (unless using heavy positional encoding techniques)
 - Usually requires numerical methods, or separate networks, to integrate




Instant Neural Graphics Primitives with a Multiresolution Hash Encoding
Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller.

Limitations of Volume Rendering

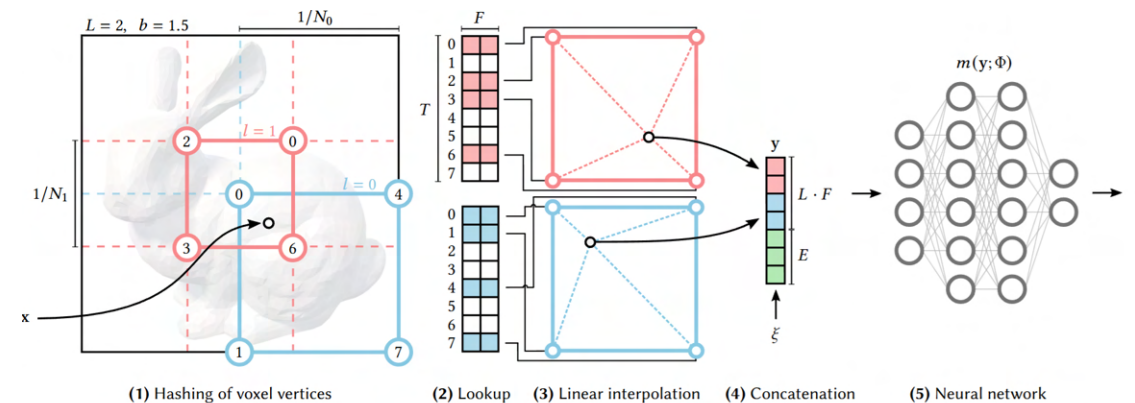
Explicit Density Fields

- Regularly spaced voxel grids, w or wo a hierarchy
 - Simple to generate (discretization) and sample from
 - Can be heavy and slow (high memory traffic)
 - Normally requires numerical methods to integrate
 - Easy to filter/downscale



Implicit Density Fields

- A neural network (MLP)
 - Difficult to generate (gradient-based optimization) and sample from (usually requires a separate model)
 - Very compact (unless using heavy positional encoding techniques)
 - Usually requires numerical methods, or separate networks, to integrate
 - Difficult to filter/downscale

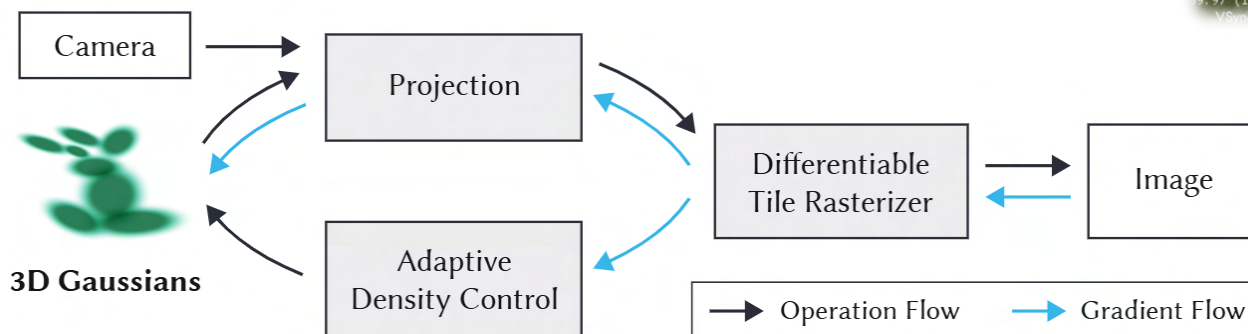


Instant Neural Graphics Primitives with a Multiresolution Hash Encoding
Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller.

Primitive-Based Volume Rendering

Defining Density Fields as Collections of Primitives

- Rasterizes collections of anisotropic Gaussians, projected in an approximate manner through EWA splatting
- Not directly applicable to PBR



3D Gaussian Splatting for Real-Time Radiance Field Rendering

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, George Drettakis

Primitive-Based Volume Rendering

Does using primitives for PBR make sense?

- They are very compact and adapt to detail (Lagrangian), scaling better than voxel grids (Eulerian)
- Potentially enables closed-form integration
- Can fit naturally to physics simulation methods (e.g. SPH fluid-solvers)



Volumetric Primitive Path Tracer (Ours)
200k Gaussians

Talk Layout

PBR with primitive volumes

Don't Splat your Gaussians: Volumetric Ray-Traced Primitives for Modeling and Rendering Scattering and Emissive Media

JORGE CONDOR, Meta Reality Labs, Switzerland and USI Lugano, Switzerland
SÉBASTIEN SPEIERER, Meta Reality Labs, Switzerland
LUKAS BODE, Meta Reality Labs, Switzerland
ALJAZ BOŽIĆ, Meta Reality Labs, Switzerland
SIMON GREEN, Meta Reality Labs, United Kingdom
PIOTR DIDYK, USI Lugano, Switzerland
ADRIÁN JARABO, Meta Reality Labs, Spain

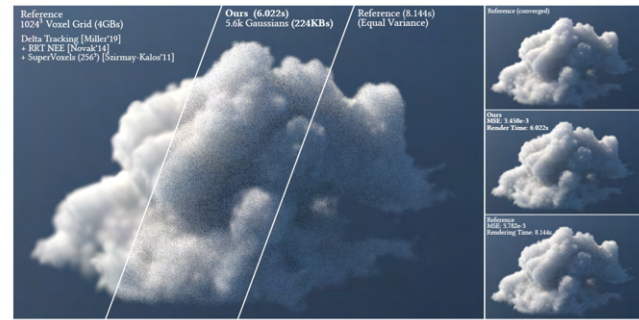


Fig. 1. We represent a complex volumetric cloud using traditional grid-based methods (left and right, 1024³ voxel grid resolution, 4GBs) and our primitives-based representation using Gaussian kernels (middle, 5.6k primitives, 224KBs), and render it with volumetric path tracing. Our method achieves substantial speedups thanks to the analytical transmission estimation and sampling, our efficient rendering approach and its extremely compact representation. When compared to the original asset, at a potential cost of detail (Figure 7), we provide large performance and memory compression benefits. Asset is part of the Walt Disney Animation Studios cloud dataset (CC-BY-SA 3.0). Rendering times reported on a NVIDIA A6000.

Authors' addresses: Jorge Condor, jorge.condor@usi.ch, Meta Reality Labs, Giesshölzstrasse 30, Zurich, Switzerland, 8045 and USI Lugano, Via Giuseppe Buffi 13, Lugano, Ticino, Switzerland, 6900; Sébastien Speierer, speierers@meta.com, Meta Reality Labs, Giesshölzstrasse 30, Zurich, Switzerland, 8045; Lukas Bode, bode@meta.com, Meta Reality Labs, Giesshölzstrasse 30, Zurich, Switzerland, 8045; Aljaz Božić, aljaz@meta.com, Meta Reality Labs, Giesshölzstrasse 30, Zurich, Switzerland, 8045; Simon Green, simongreen@meta.com, Meta Reality Labs, London, United Kingdom, 8045; Piotr Didyk, piotr.didyk@usi.ch, USI Lugano, Via Giuseppe Buffi 13, Lugano, Switzerland, 6900; Adrián Jarabo, ajarabo@meta.com, Meta Reality Labs, Zaragoza, Spain.

Please use nonem option or ACM Engage class to enable CC licenses. This work is licensed under a Creative Commons Attribution 4.0 International License. ACM 978-0-9959-2025-2/ART10
<https://doi.org/10.1145/3711853>

ACM ToG'24 (presented at SIGGRAPH'25)

LOD for PBR with primitive volumes

Gabor Fields: Orientation-Selective Level-of-Detail for Volume Rendering

JORGE CONDOR*, USI Lugano, Switzerland
NICOLAI HERMANN*, USI Lugano, Switzerland
MEHMET ATA YURTSEVER, USI Lugano, Switzerland
PIOTR DIDYK, USI Lugano, Switzerland



Fig. 1. Example level-of-detail decomposition of our Gabor Fields, a new volumetric density field representation using primitive Gabor kernels. The frequency and orientation-selectivity of the representation enables continuous control over level of detail, reduced rendering time, and a trade-off between variance and bias for performance when sample cost is critical. Here, we show four levels obtained by masking particle collections from the highest-quality asset at render time. (Volumetric Blurry from the OpenVDB database [DuramWorkshopAnimation 2023])

Gaussian-based representations have enabled efficient physically-based volume rendering at a fraction of the memory cost of regular, discrete, voxel-based distributions. However, several remaining issues hamper their widespread use. One of the advantages of classic voxel grids is the ease of constructing hierarchical representations by either storing volumetric mipmaps or selectively pruning branches of an already hierarchical voxel grid. Such strategies reduce rendering time and eliminate aliasing when lower levels of detail are required. Constructing similar strategies for Gaussian-based volumes is not trivial. Straightforward solutions, such as prefiltering or computing mipmap-style representations, lead to increased memory requirements or expensive re-fitting of each level separately. Additionally, such solutions do not guarantee a smooth transition between different hierarchy levels. To address these limitations, we propose Gabor Fields, an orientation-selective mixture of Gabor kernels that enables continuous frequency filtering at no cost. The frequency content of the asset is reduced by selectively pruning primitives, directly benefiting rendering performance. Beyond filtering, we demonstrate that stochastic sampling from different frequencies and orientations at each ray recursion enables masking substantial portions of the volume, accelerating ray traversal time in single- and multiple-scattering settings. Furthermore, inspired by procedural volumes,

*Both authors contributed equally to this research.

Authors' Contact Information: Jorge Condor, jorge.condor@usi.ch, USI Lugano, Lugano, Ticino, Switzerland; Nicolai Hermann, nicolai.hermann@usi.ch, USI Lugano, Lugano, Switzerland; Mehmet Ata Yurtsever, ata.yurtsever@usi.ch, USI Lugano, Lugano, Switzerland; Piotr Didyk, piotr.didyk@usi.ch, USI Lugano, Lugano, Switzerland.

ACM ToG (SIGGRAPH'26) (conditionally accepted)

Neural Primitive Volumes

Neural Harmonic Textures for High-Quality Primitive Based Neural Reconstruction

Jorge Condor^{1,2}, Nicolas Moëne-Loccoz¹, Merlin Nimier-David¹, Piotr Didyk², Zan Gojic¹, and Qi Wu¹

¹ NVIDIA
{nicolasm,mmierdavid,zgojcic,qiwu}@nvidia.com
² Università della Svizzera italiana, Lugano, Switzerland
{jorge.condor,piotr.didyk}@usi.ch




Fig. 1: Neural Harmonic Textures for novel view synthesis. We attach learnable feature vectors (right) to the virtual vertices of bounding tetrahedra encapsulating each primitive (center). After harmonic encoding and accumulation along the ray, a small neural network decodes the resulting signal into RGB color in a deferred manner (left). Source code and further results are available at <https://research.nvidia.com/labs/s11/projects/neural-harmonic-textures/>.

Abstract. Primitive-based methods such as 3D Gaussian Splatting have recently become the state-of-the-art for novel-view synthesis and related reconstruction tasks. Compared to neural fields, these representations are more flexible, adaptive, and scale better to large scenes. However, the limited expressivity of individual primitives makes modeling high-frequency detail challenging. We introduce *Neural Harmonic Textures*, a neural representation approach that anchors latent feature vectors on a virtual scaffold surrounding each primitive. These features are interpolated within the primitive at ray intersection points. Inspired by Fourier analysis, we apply periodic activations to the interpolated features, turning alpha blending into a weighted sum of harmonic components. The resulting signal is then decoded in a single deferred pass using a small

Arxiv'26 (under revision)



Don't Splat your Gaussians: Volumetric Ray Traced Primitives for Modelling and Rendering Scattering and Emissive Media

Jorge Condor

Sébastien Speierer


Lukas Bode

Aljaz Bozic


Simon Green

Piotr Didyk

Adrián Jarabo



Università
della
Svizzera
italiana



Recipe for a PBR Volume Rendering Method

Ingredients

- A density field
- A way to regress/obtain this field
- A way to integrate density along rays slicing this field
- A way to sample distances (aka how to scatter)
- An efficient implementation of all the above within a path tracer

Path Tracing Primitive Volumes

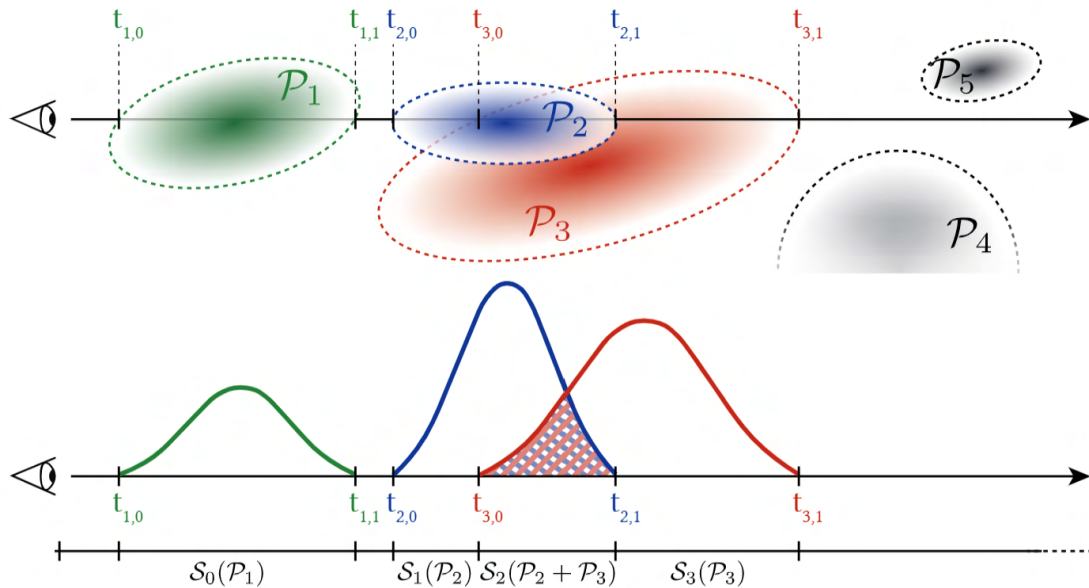
Density Field as Mixture of Primitives



Path Tracing Primitive Volumes


Analytic Transmittance Integrals

- Closed-form solutions to the line integrals of the anisotropic kernels along the ray
- Bounds define edges, we integrate segment by segment analytically



Path Tracing Primitive Volumes

Distance Sampling in Primitive Volumes



Path Tracing Primitive Volumes

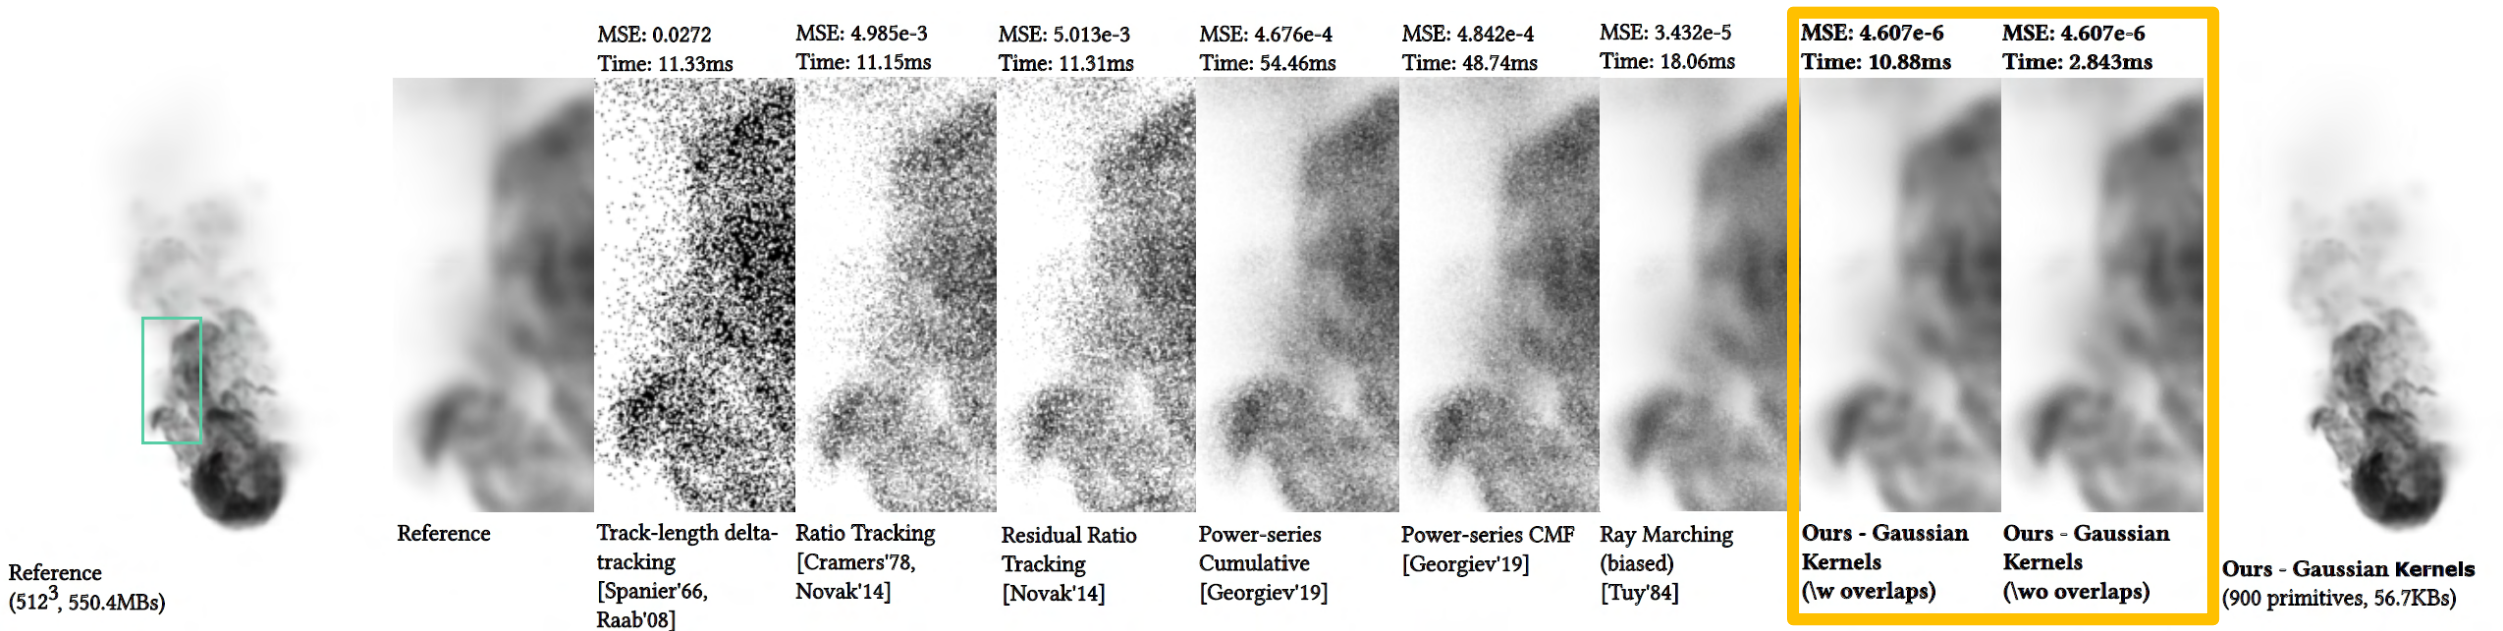
Regressing Assets as Mixtures of Primitives

- Adjoint renderer based on path-replay backpropagation
- Inverse renderer can optimize absorbing and scattering volumes from images




Results

Tomography



Results

Unbiased Multiple Scattering



Results

Unbiased Multiple Scattering

Higher Quality, Slower

Less Primitives, Faster



VDB (480MB)
(Reference)

16.2k (639KB)
MSE: 1.621e-4
26.27mins

5.6k (224KB)
MSE: 1.717e-4
19.43mins

2.6k (107KB)
MSE: 2.72e-4
11.67mins

Results



Results

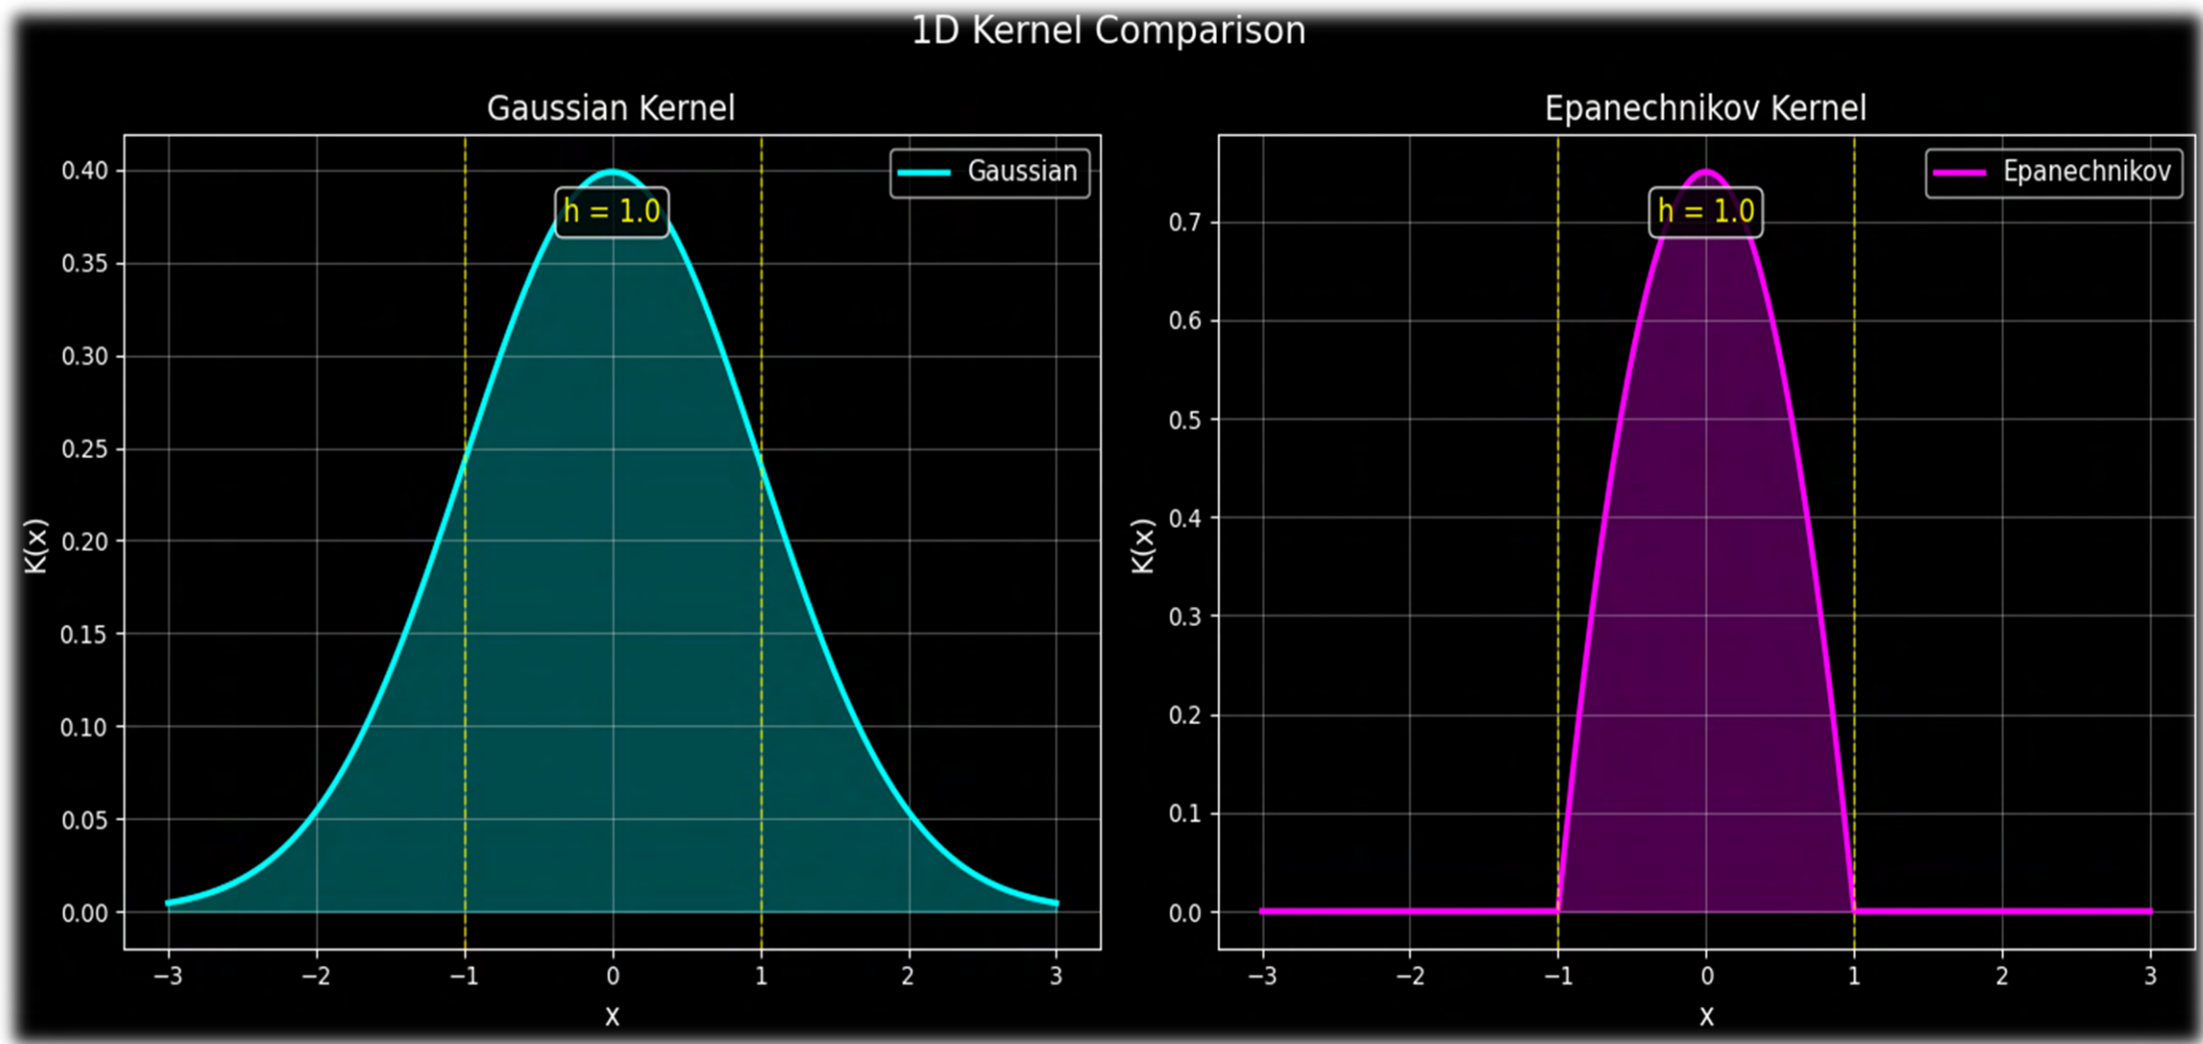
Radiance Fields with Complex Cameras



Limitations

- Optimization can be slow (0.5-2h)
- Performance depends on measures like amount of overlapping primitives
→ can become slow with too many primitives.
- Level of detail?

A Kernel-Agnostic Formulation



A Kernel-Agnostic Formulation



Epanechnikov: 7.653 ms



Gaussian: 24.966 ms

Gabor Fields: Orientation-Selective Level-of-Detail for Volume Rendering

Jorge Condor*


Nicolai Hermann*

Mehmet Ata Yurtsever

Piotr Didyk



Università
della
Svizzera
italiana




Constructing Level of Detail with Gaussian Primitives

Why is LOD challenging with primitives

- Analytic filtering is possible
- In practice: very slow to render!
 - Same number of primitives
 - Larger (higher overlap)
 - Higher memory footprint (need to store each level independently)


Analytic Gaussian
Filtering



Constructing Level of Detail with Gaussian Primitives


Alternatives

- Retrain the same asset at different levels of detail (pyramid) → slow, extra memory, requires manual tuning
- Heuristic pruning of small Gaussians




What is the core issue?

- Gaussians in the Frequency domain span all frequencies



What is the core issue?

- Gaussians in the Frequency domain span all frequencies



LOD for Primitive Volumes

What do we want to achieve?

- We want to train once, get directly an asset that can we can filter arbitrarily without post processing, in real time
- No extra memory
- Efficient to evaluate and integrate

LOD for Primitive Volumes

What do we want to achieve?

- We want to train once, get directly an asset that can we can filter arbitrarily without post processing, in real time
- No extra memory
- Efficient to evaluate and integrate

Anisotropic Gaussian Envelope


$$e^{-\frac{1}{2} (x-\mu)^T \Sigma^{-1} (x-\mu)} \cos \left(\vec{\omega}^T (x - \mu) \right)$$

Harmonic Modulation




Kernels as Frequency Residuals

- Gabor kernels are defined in specific frequency bands \rightarrow filtering reduces to deleting part of the volume

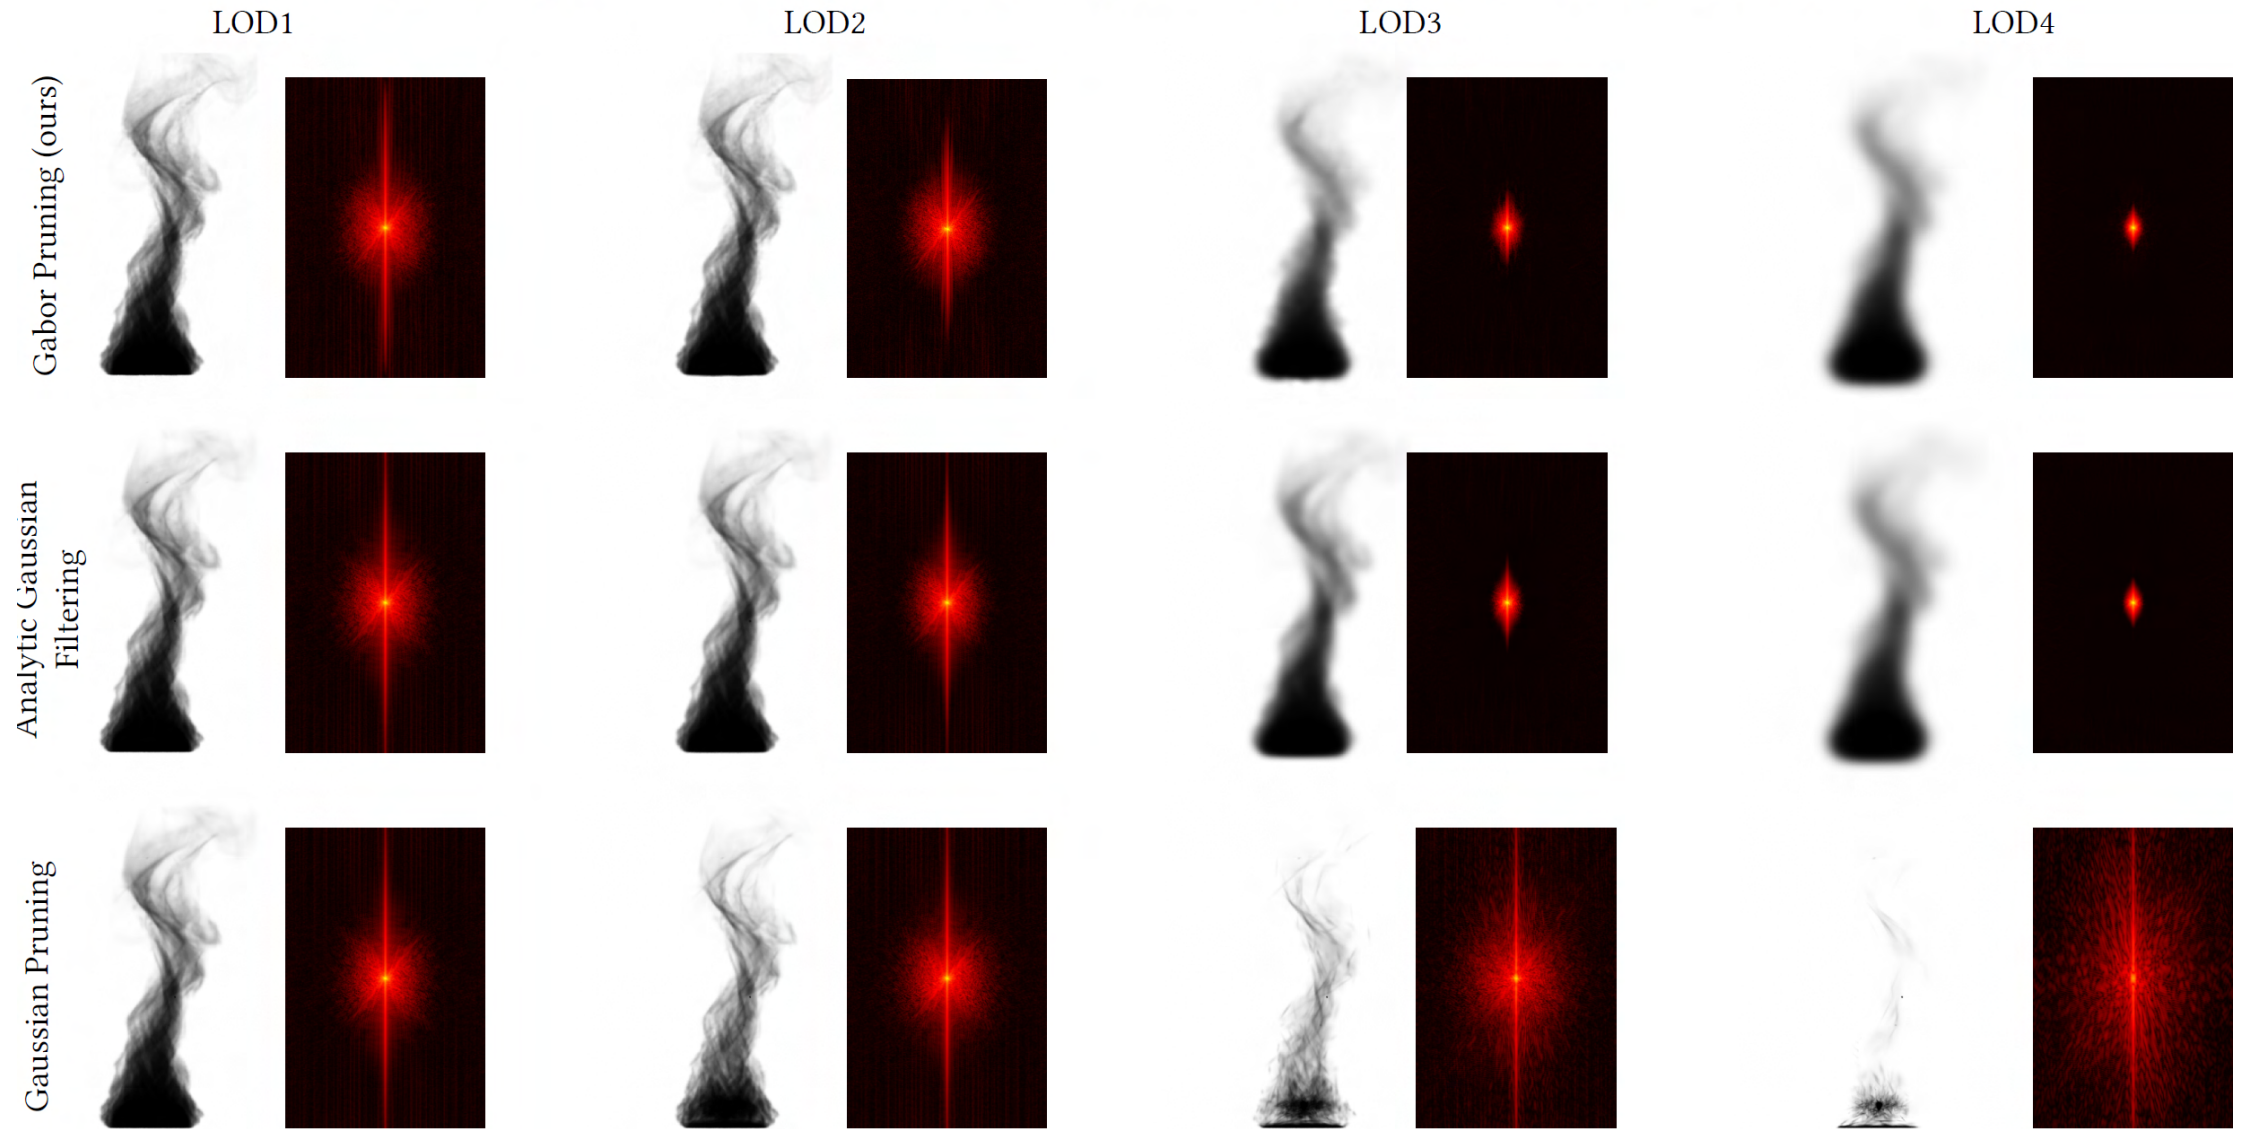


Kernels as Frequency Residuals

- Gabor kernels are defined in specific frequency bands \rightarrow filtering reduces to deleting part of the volume



LOD Decomposition Methods: Rendered Images and Log Power Spectra (Frequency Domain)



Gabor Fields


Contributions

- Defining volumes as mixtures of Gaussians (zero-freq. Gabors) and Gabor kernels
- Deriving closed-form line integrals for the full domain and segments
- Analytic Fourier Transforms to filter arbitrarily and continuously
- Stochastic-analytic methods to trade sample cost and variance, accelerating MS
- Implemented using Optix ray masks to accelerate BBVH traversal

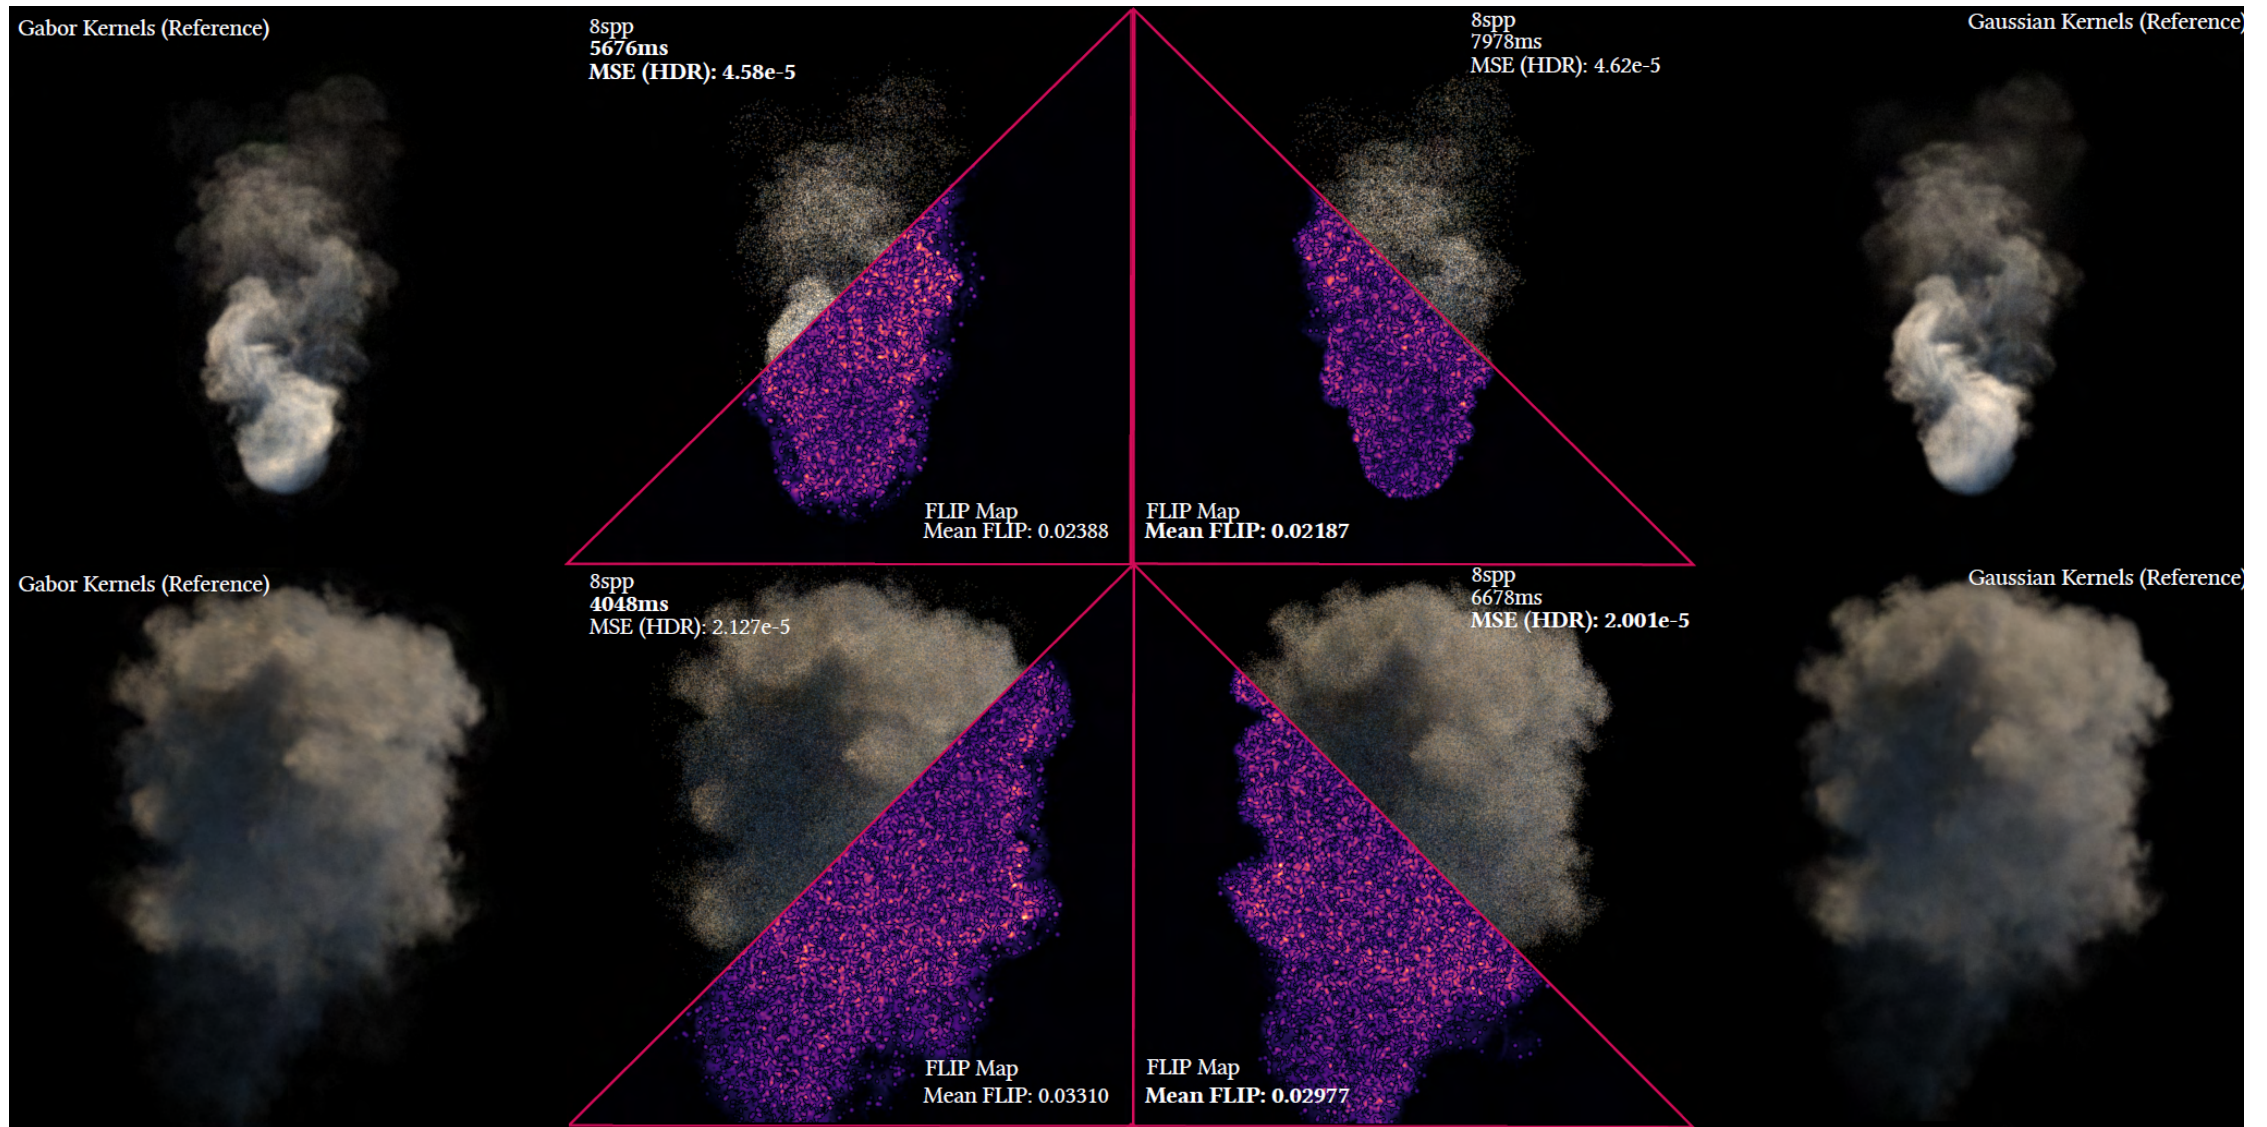


Tornado, 32k Gabor kernels

Gabor Fields




Gabor Fields



Gabor Fields: LOD

LOD



LoD Rendering
Frame: 0

PRIMITIVES
33,280

LEVELS
8 / 8

EXEC TIME
584.9 ms

DISTANCE
5.00

LOD (static)



LoD Rendering
Frame: 0

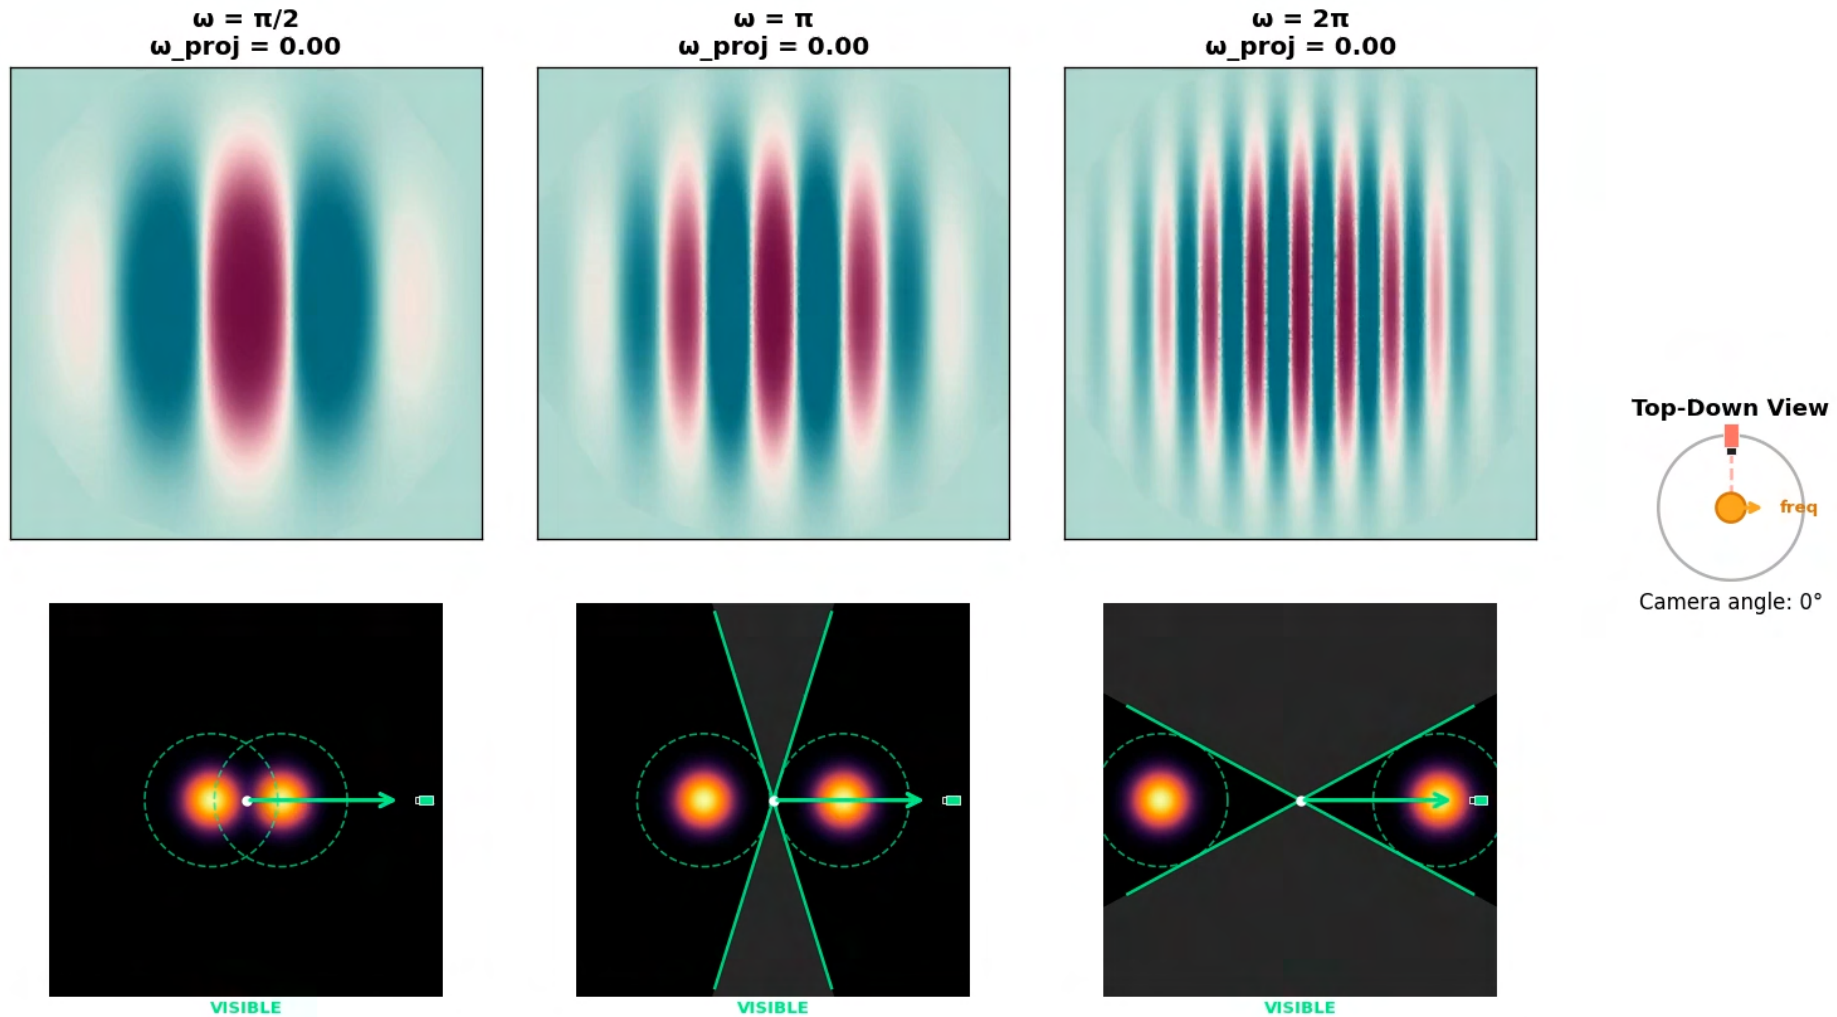
PRIMITIVES
33,280

LEVELS
8 / 8

EXEC TIME
626.4 ms



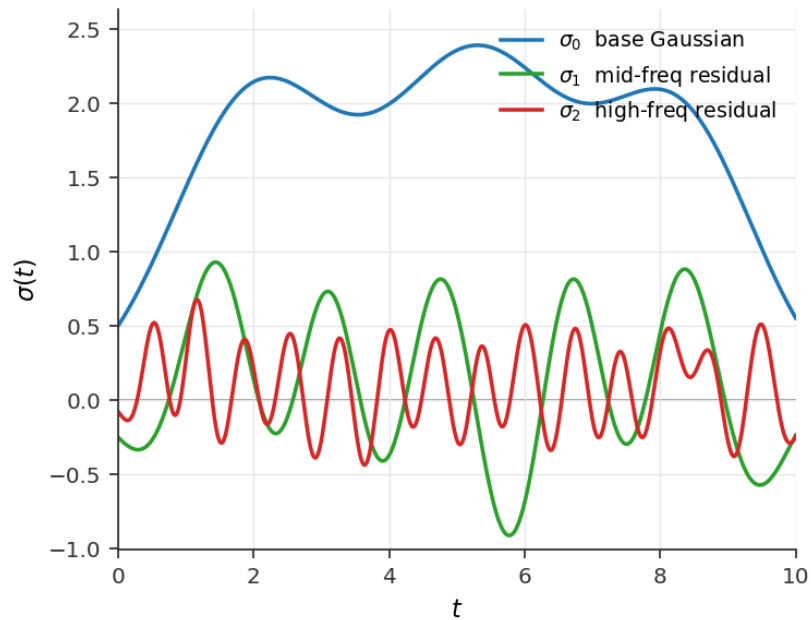
DISTANCE
5.00

Frequency and Orientation Selectivity in Gabor Fields






Stochastic-Analytic Estimators

Can we use the frequency properties of Gabor kernels to further reduce sample cost?



Stochastic-Analytic Estimators

Sampling from frequency levels





Stochastic-Analytic Tomography with Gabor Fields

Reference

Uniform Sampling (1spp)

Control Variate
Power Law (1spp)



- Deterministic (Reference)
- Uniform Stochastic
- Power Law
- Control Variate Uniform
- Control Variate Power Law



Stochastic-Analytic Path Tracing with Gabor Fields

Ours (unbiased)



512spp
Execution time: 490.26s

Ours
+ Control Variate PLA




512spp
Execution time: 295.47s

Ours
+ Control Variate PLA
+ Zero NEE

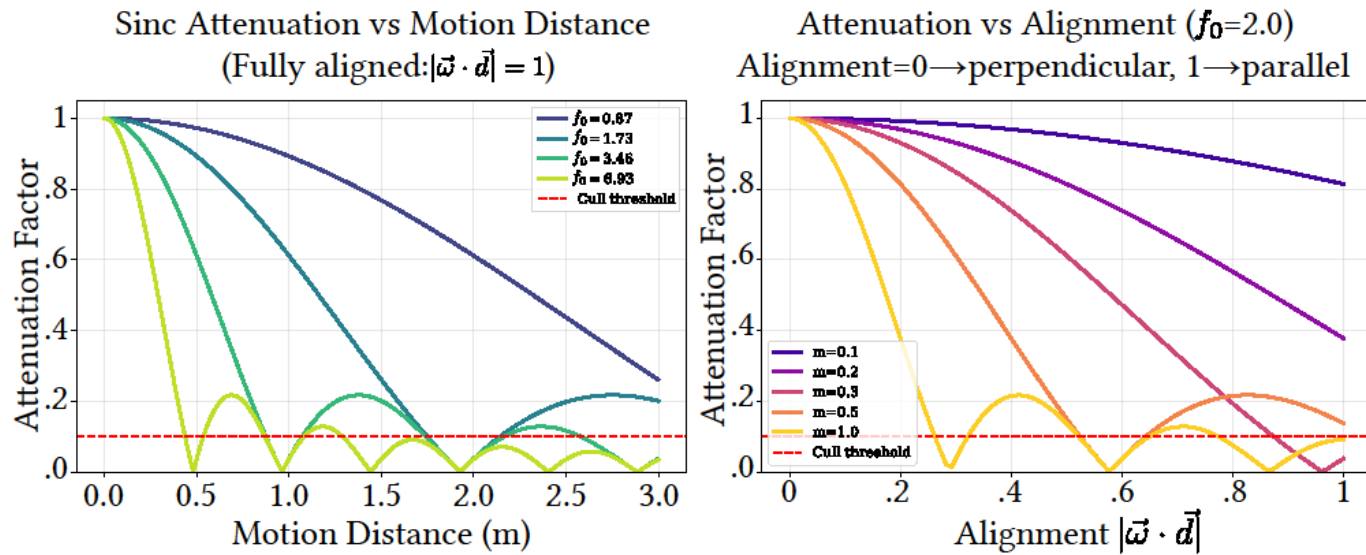


512spp
Execution time: 239.95s

Foveated Rendering with Gabor Fields




Efficient Motion Blur with Gabor Fields



Limitations, Future Work

- Primitive volumes are still expensive to render for many overlapping primitives

- Gabor kernels can be loosely interpreted as the harmonics composing local density field signals
- Can we lift this to higher dimensions and use it for scene reconstruction?



Neural Harmonic Textures for High-Quality Primitive-Based Reconstruction

Jorge Condor

Nicolas Moënné-Loccoz

Merlin Nimier-David

Piotr Didyk

Zan Gojcic


Qi Wu



Università
della
Svizzera
italiana




NVIDIA®



Arxiv'26 (under submission)


The appearance modeling problem

- 3DGS models appearance and geometry jointly --> this limits fundamentally scalability and detail modelling



Textured Gaussians for Enhanced 3D Scene Appearance Modeling
Chao et. al, CVPR'25

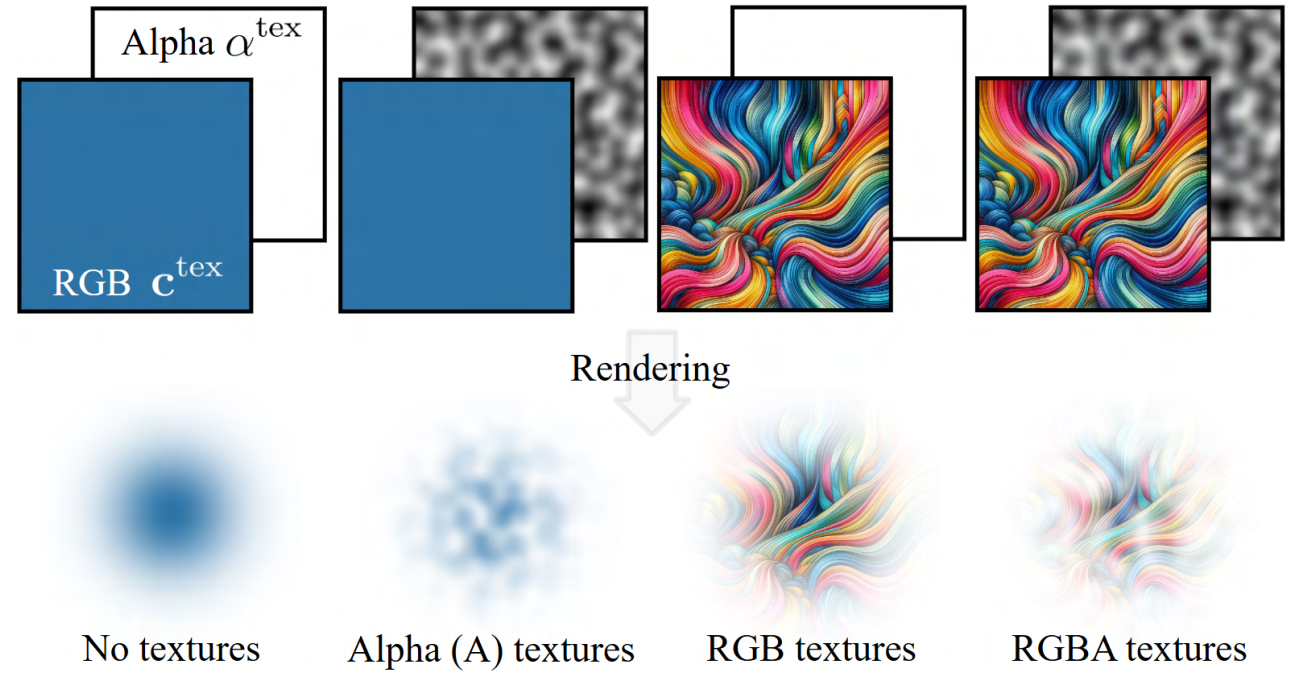
- Limited expressivity of individual Gaussians + frequency-limited SHs



Spherical Harmonics

Towards more expressive appearance models

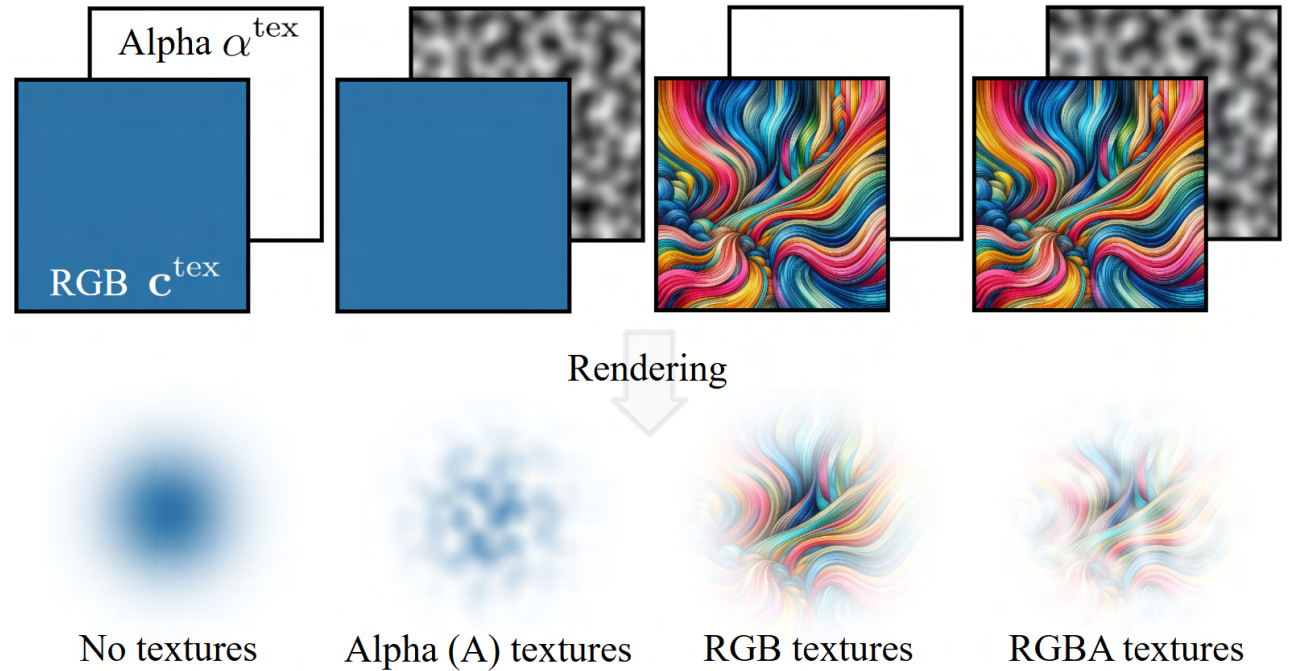
- Previous work:
 - Texturing gaussians
 - Parametric spherical functions (spherical Betas, Voronoi)
 - 3DGS geometry + neural field appearance (radiance meshes, neural shell textures)



Textured Gaussians for Enhanced 3D Scene Appearance Modeling
Chao et. al, CVPR'25

Towards more expressive appearance models

- Previous work:
 - Texturing gaussians
 - Parametric spherical functions (spherical Betas, Voronoi)
 - 3DGS geometry + neural field appearance (radiance meshes, neural shell textures)
- Issues:
 - Complex optimization landscapes
 - Expensive evaluation (heavy textures, excessive MLP queries)
 - Complete detachment of neural features from geometry (geometry scales and moves, but the field does not adapt a it has its own structure)




Textured Gaussians for Enhanced 3D Scene Appearance Modeling
Chao et. al, CVPR'25

Revisiting Neural Radiance Fields

Neural Fields can be the solution to these problems!

- Can pool information across the scene (if using global fields)
- Detach the geometry/3D encoding structure from the appearance
- More expressive
 - They can model signals of arbitrary frequency (no freq. banding)



Instant Neural graphics Primitives with a Multiresolution Hash Encoding. Müller et al. SIGGRAPH'22

Lifting Primitive Appearance Through Neural Fields

Can we preserve the benefits of volumetric primitives?

- A) Each primitive is its own Neural Field
- B) Primitives are geometry, a globally-supported Neural Field handles appearance



Splat the Net: Radiance Fields with Splattable Neural Primitives. Zhou et. al. ICLR'26

Radiance Meshes for Volumetric Reconstruction. Mai. et. al. CVPR'26

Lifting Primitives Through Neural Fields

Can we preserve the benefits of volumetric primitives?

- A) Each primitive is its own Neural Field
- B) Primitives are geometry, a globally-supported Neural Field handles appearance



Splat the Net: Radiance Fields with Splattable Neural Primitives. Zhou et. al. ICLR'26


Radiance Meshes for Volumetric Reconstruction. Mai. et. al. CVPR'26

- Very complex optimization landscape
- Expensive closed integration

- Bad scalability
- Subpar quality (related to scalability)
- Expensive (dozens of MLP queries per pixel)
- Cannot easily support motion without baking

Neural Harmonic Textures: Overview


- 1) 3D Gaussians as both geometry and local fields
- 2) Harmonic Texturing
- 3) Neural Deferred Shading



Neural Harmonic Textures

Primitives as local fields


- Virtual tetrahedron defines a local field
- Each vertex is an N-dimensional feature vector
- Response obtained through interpolation at evaluation point



Neural Harmonic Textures: Overview

Neural Deferred Shading with a Shallow MLP


- Efficient implementation in regular engines
- Very fast (single query per pixel)
- Simpler optimization landscape



Neural Harmonic Textures

Harmonic Textures


- Encoding response with a harmonic function
 - Analogous to accumulating the harmonics of the signal in Fourier decompositions



Neural Harmonic Textures

Harmonic Textures

- Encoding response with a harmonic function
 - Analogous to accumulating the harmonics of the signal




Results

Method	MipNeRF360			Tanks & Temples			Deep Blending		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
Instant NGP-Big	25.59	0.695	0.375	21.92	0.740	0.342	24.96	0.815	0.459
Mip-NeRF 360	27.60	0.788	0.275	22.22	0.754	0.290	29.40	0.899	0.306
ZipNeRF	28.55	0.829	0.218	23.64	0.836	0.179	—	—	—
2DGS	27.22	0.804	0.275	22.85	0.827	0.244	29.56	0.904	0.325
3DGS-MCMC	27.99	0.830	0.229	24.46	0.866	0.174	29.49	0.912	0.306
3DGUT-MCMC	27.82	0.826	0.233	24.20	0.861	0.180	29.87	0.913	0.309
Beta Splatting-MCMC	28.12	0.831	0.238	24.54	0.866	0.196	29.56	0.907	0.316
Spherical Voronoi	28.56	0.835	0.228	24.80	0.871	0.172	30.34	0.914	0.299
Triangle Splatting	27.00	0.808	0.231	23.05	0.843	0.191	28.92	0.891	0.308
Textured Gaussians	27.35	0.827	—	24.26	0.854	—	28.33	0.891	—
NeST	26.54	0.776	0.260	—	—	—	—	—	—
Radiance Meshes	27.15	0.810	0.274	23.13	0.851	0.200	29.39	0.901	0.362
Neural Harmonic Textures (Ours)	28.74	0.834	0.216	25.68	0.882	0.141	30.94	0.919	0.302



Results



3DGUT-MCMC

Spherical Voronoi

Ours

Ground truth



Results: 10k Primitives



Reference

3DGUT-MCMC

Ours

PSNR: 23.97dB

SSIM: 0.724

LPIPS (VGG): 0.548

PSNR: **27.63dB** (+3.7)

SSIM: **0.856** (+0.132)

LPIPS (VGG): **0.410** (-0.138)

Results: 10k Primitives



Reference

3DGUT-MCMC

Ours

PSNR: 23.97dB

PSNR: **27.63dB** (+3.7)


SSIM: 0.724

SSIM: **0.856** (+0.132)

LPIPS (VGG): 0.548

LPIPS (VGG): **0.410** (-0.138)


Results: Scalability



Result: Efficient Arbitrary Signal Fitting





Real-time Intrinsic Reconstruction (PBR)



Semantic Reconstruction
(512-dimensional vectors)

Results: 2D HDR Image Regression

Alternative formulation using NHT as a Lagrangian Positional Encoding → Single Opaque, Connected 2D Mesh



Results: 2D HDR Image Regression

Alternative formulation using NHT as a Lagrangian Positional Encoding



Reference (45MP, 14-bit HDR)

Instant NGP (100x)

Ours (100x)

Limitations

- Like all expressive appearance models in scene reconstruction: prone to overfitting in sparse captures
- Slower than 3DGS

Method (w/ MCMC)	MipNeRF360				Tanks & Temples				Deep Blending			
	PSNR↑	SSIM↑	LPIPS↓	FPS↑	PSNR↑	SSIM↑	LPIPS↓	FPS↑	PSNR↑	SSIM↑	LPIPS↓	FPS↑
3DGS + SH	27.94	0.829	0.246	251	24.25	0.861	0.188	294	29.98	0.912	0.317	331
3DGUT + SH	27.93	0.828	0.247	201	23.99	0.859	0.192	245	30.21	0.913	0.318	282
3DGUT + SV	28.15	0.823	0.248	202	24.18	0.861	0.187	242	30.29	0.912	0.320	267
3DGUT + NHT (Ours)	28.46	0.830	0.232	140	24.79	0.875	0.169	226	30.88	0.918	0.311	240

1M primitives, all methods implemented in *gsplat*

Future Work

- Other applications: radiance caching for GI, relighting, downstream applications from having real-time LSEG/DINO features.
- Gradients can now propagate to geometric features from the appearance directly
 - Remove Gaussian kernels
 - Use cheaper primitives (e.g. tetrahedrons)

Reach out!



Physically-Based Volume Rendering and High-quality Neural Reconstruction with Volumetric Primitives

Jorge Condor
PhD Student @ USI Lugano, Switzerland

21st April 2026
Saarbrücken, Germany