Machine Learning in Rendering

Part-4

Gurprit Singh / Philipp Slusallek / Karol Myszkowski

Overview

Radiance regression functions for global illumination

Contrastive learning in path manifolds

Neural radiance Fields

Neural control variates

Non-exponential transmittance model for volumetric scene representations

Compositional neural scene representations for shading inference

Radiance regression functions

Pre-computed Radiance Transfer (PRT) methods

Precomputes the global light transport

Stores the resulting PRT for fast rendering

Challenging scenarios: dynamic viewpoint and lighting

- e.g., dynamic local lights & glossy interreflections

Real-time PRT renderings

Glossy inter-reflections & dynamic lights

Radiance Regression Functions

A function that returns the indirect illumination value for each surface point given the viewing direction and lighting condition.

RRF is learned using non-linear regression

Dataset: training samples precomputed by offline rendering

Results in a real-time rendering of precomputed global illumination

Radiance Regression Functions

First component

Basic attributes: surface position (2D), location of light source (2D) and the viewing direction (2D)

— not sufficient to capture spatially-variant surface properties

Augmented attributes: surface normals and material properties

RRF directly approximates global illumination: highly complex and non-linear 6D function

Existing PRT methods only exploits nonlinear coherence in some dimensions and perform dense sampling in the other dimensions.

Radiance Regression Functions

Second component

Partition the space and fit a separate RRF for each of the subspaces

Multiple small MLPs that collectively and efficiently represent indirect illumination

Radiance Regression Functions

Full pipeline

Render caustics, sharp indirect shadows, high-frequency glossy reflections

The precomputed network only depends on surface and not the underlying surface meshing

Makes it more scalable than PRT methods

Shows 30FPS with 512x512 images

Efficiency depends on screen size and not the scene

Radiance Regression Functions

MLP structure

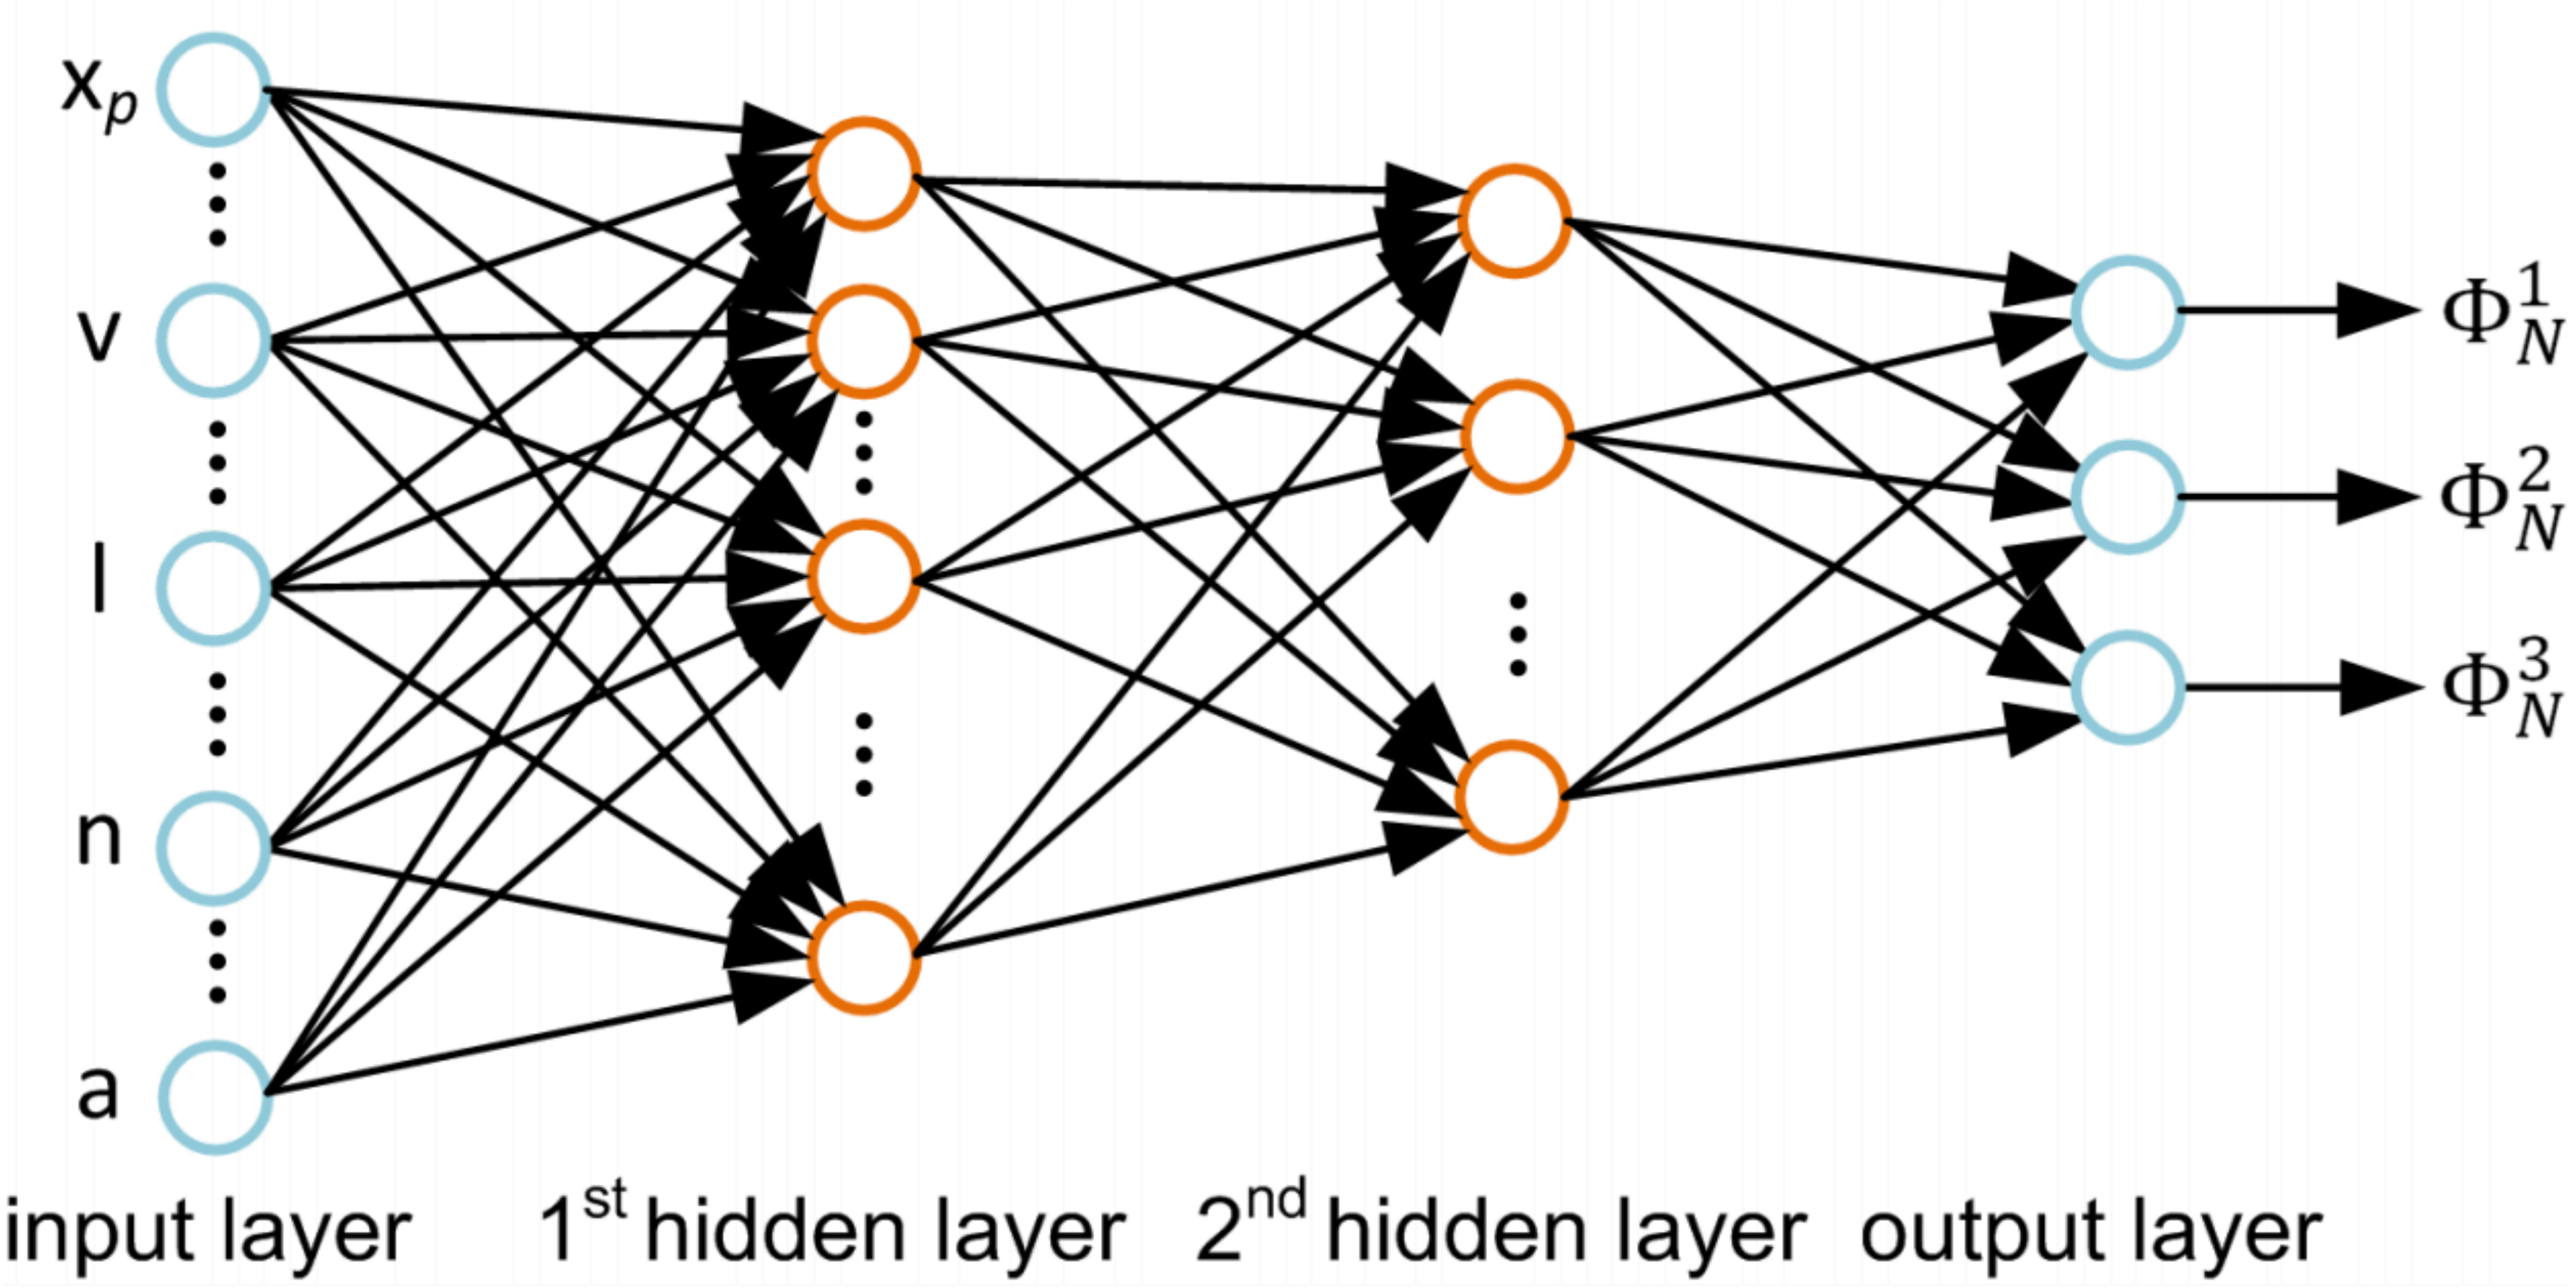

Radiance Regression Functions

MLP structure

Activation function: hyperbolic tangent function

$$tanh(z) = \frac{2}{1 - e^{-2z}} - 1$$

Radiance Regression Functions

Results

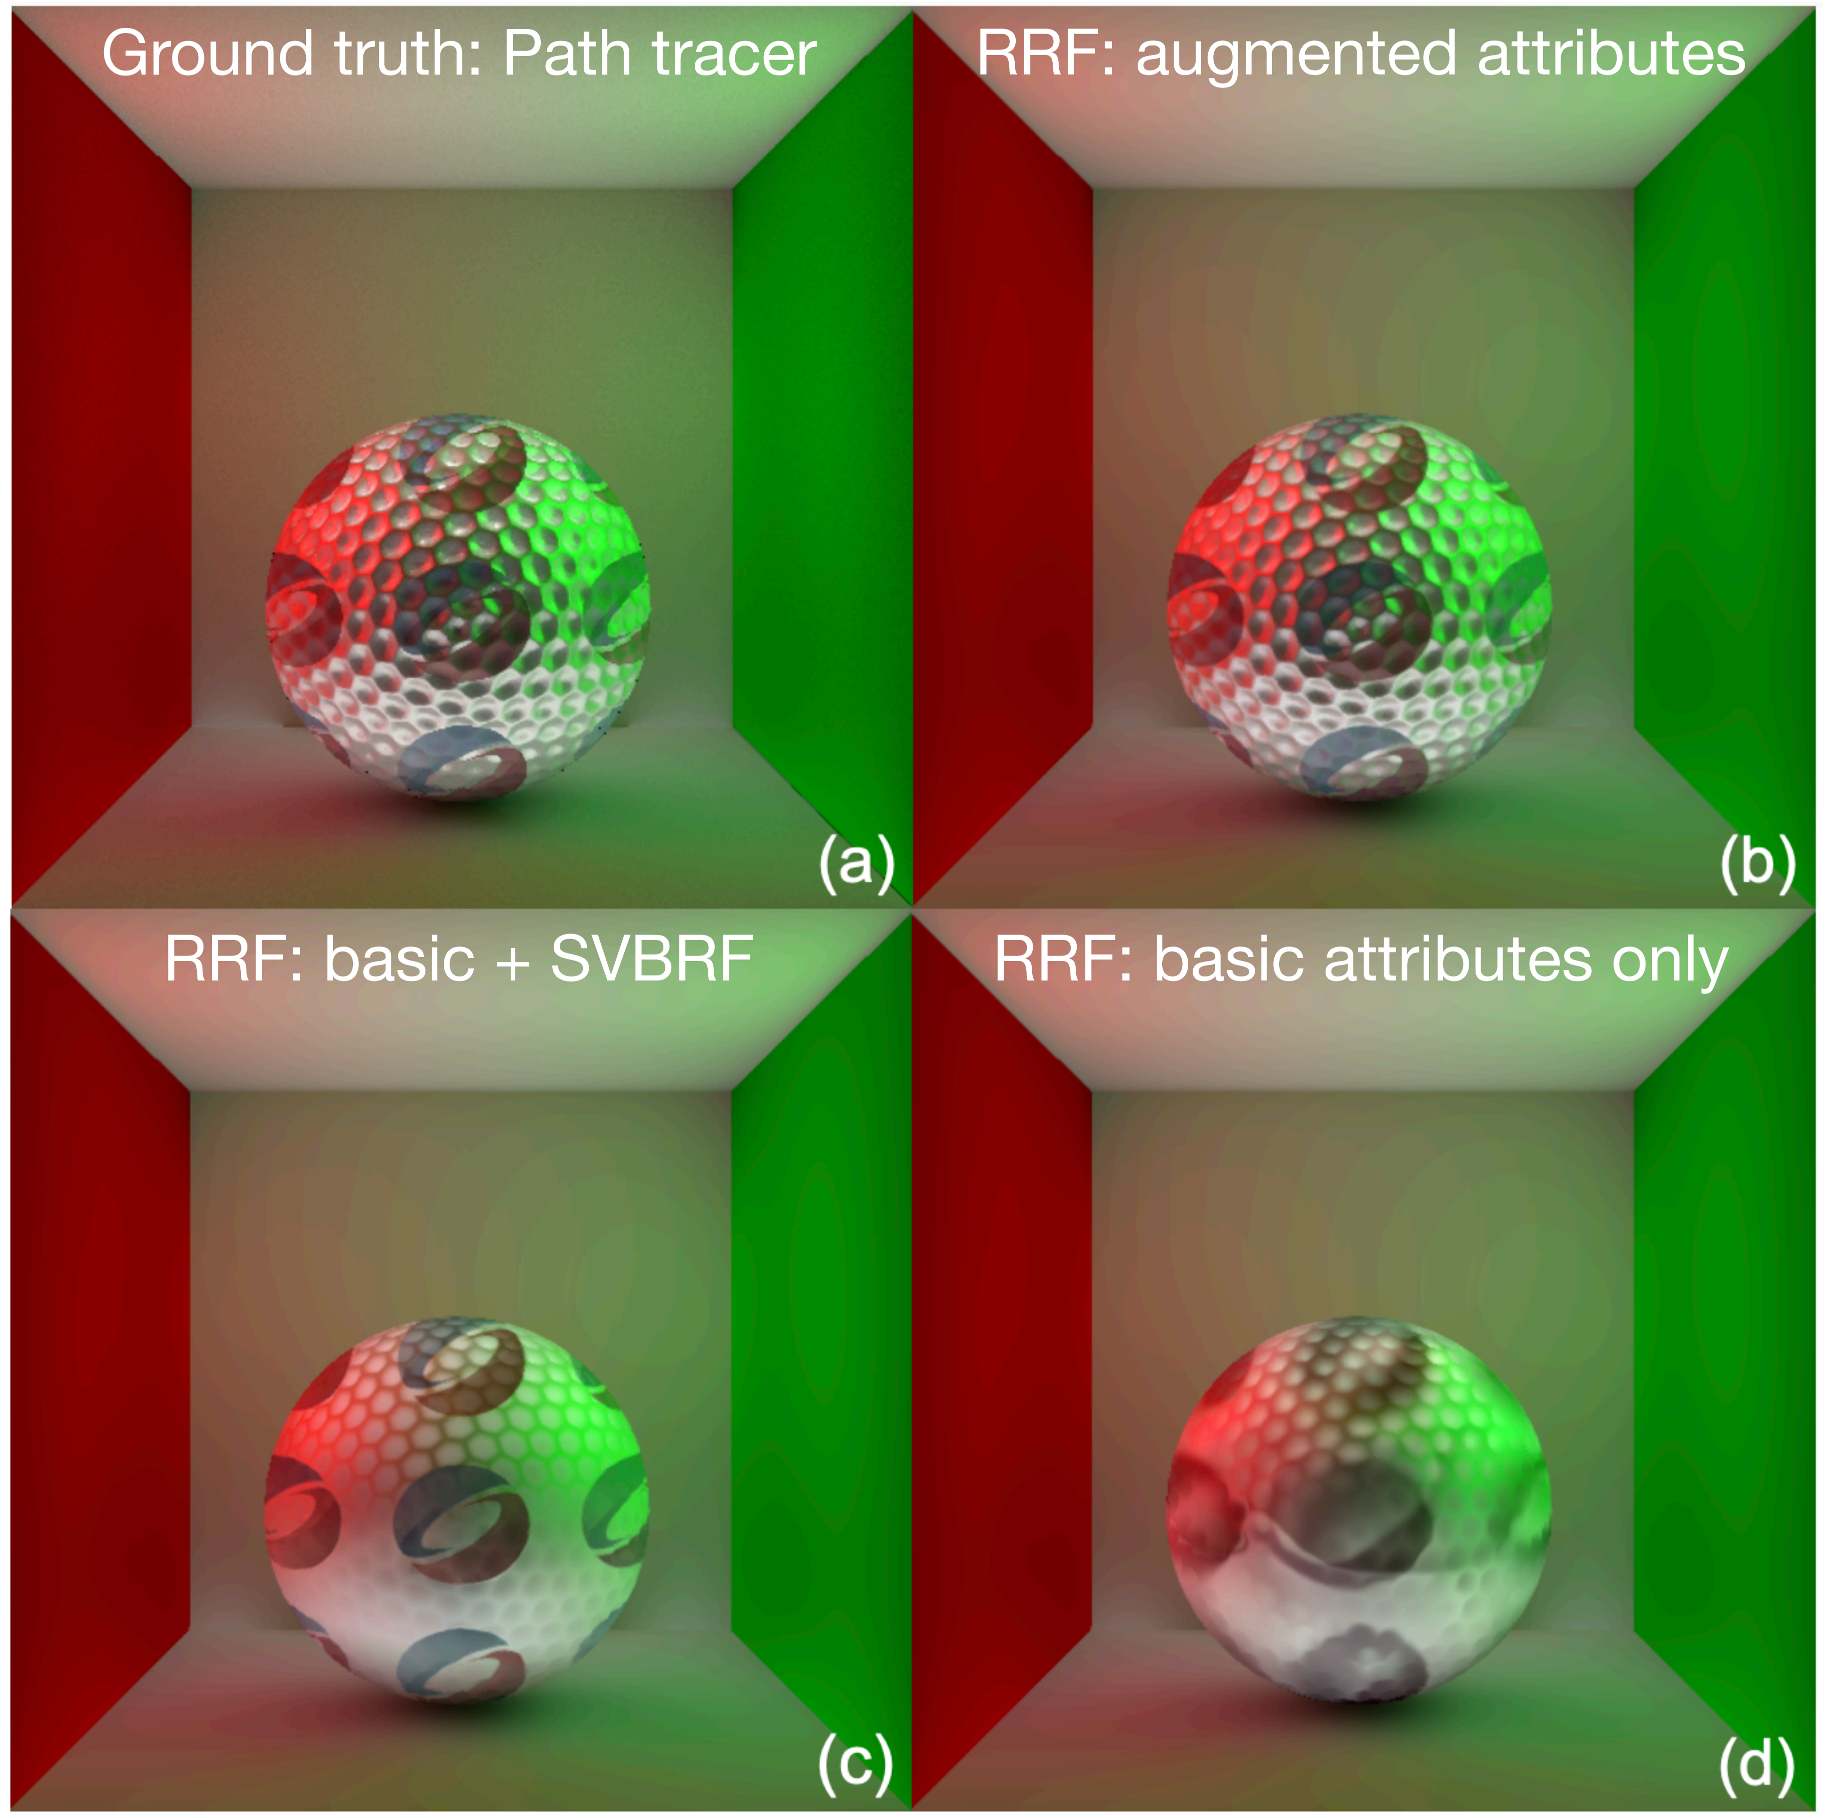

Radiance Regression Functions

Results

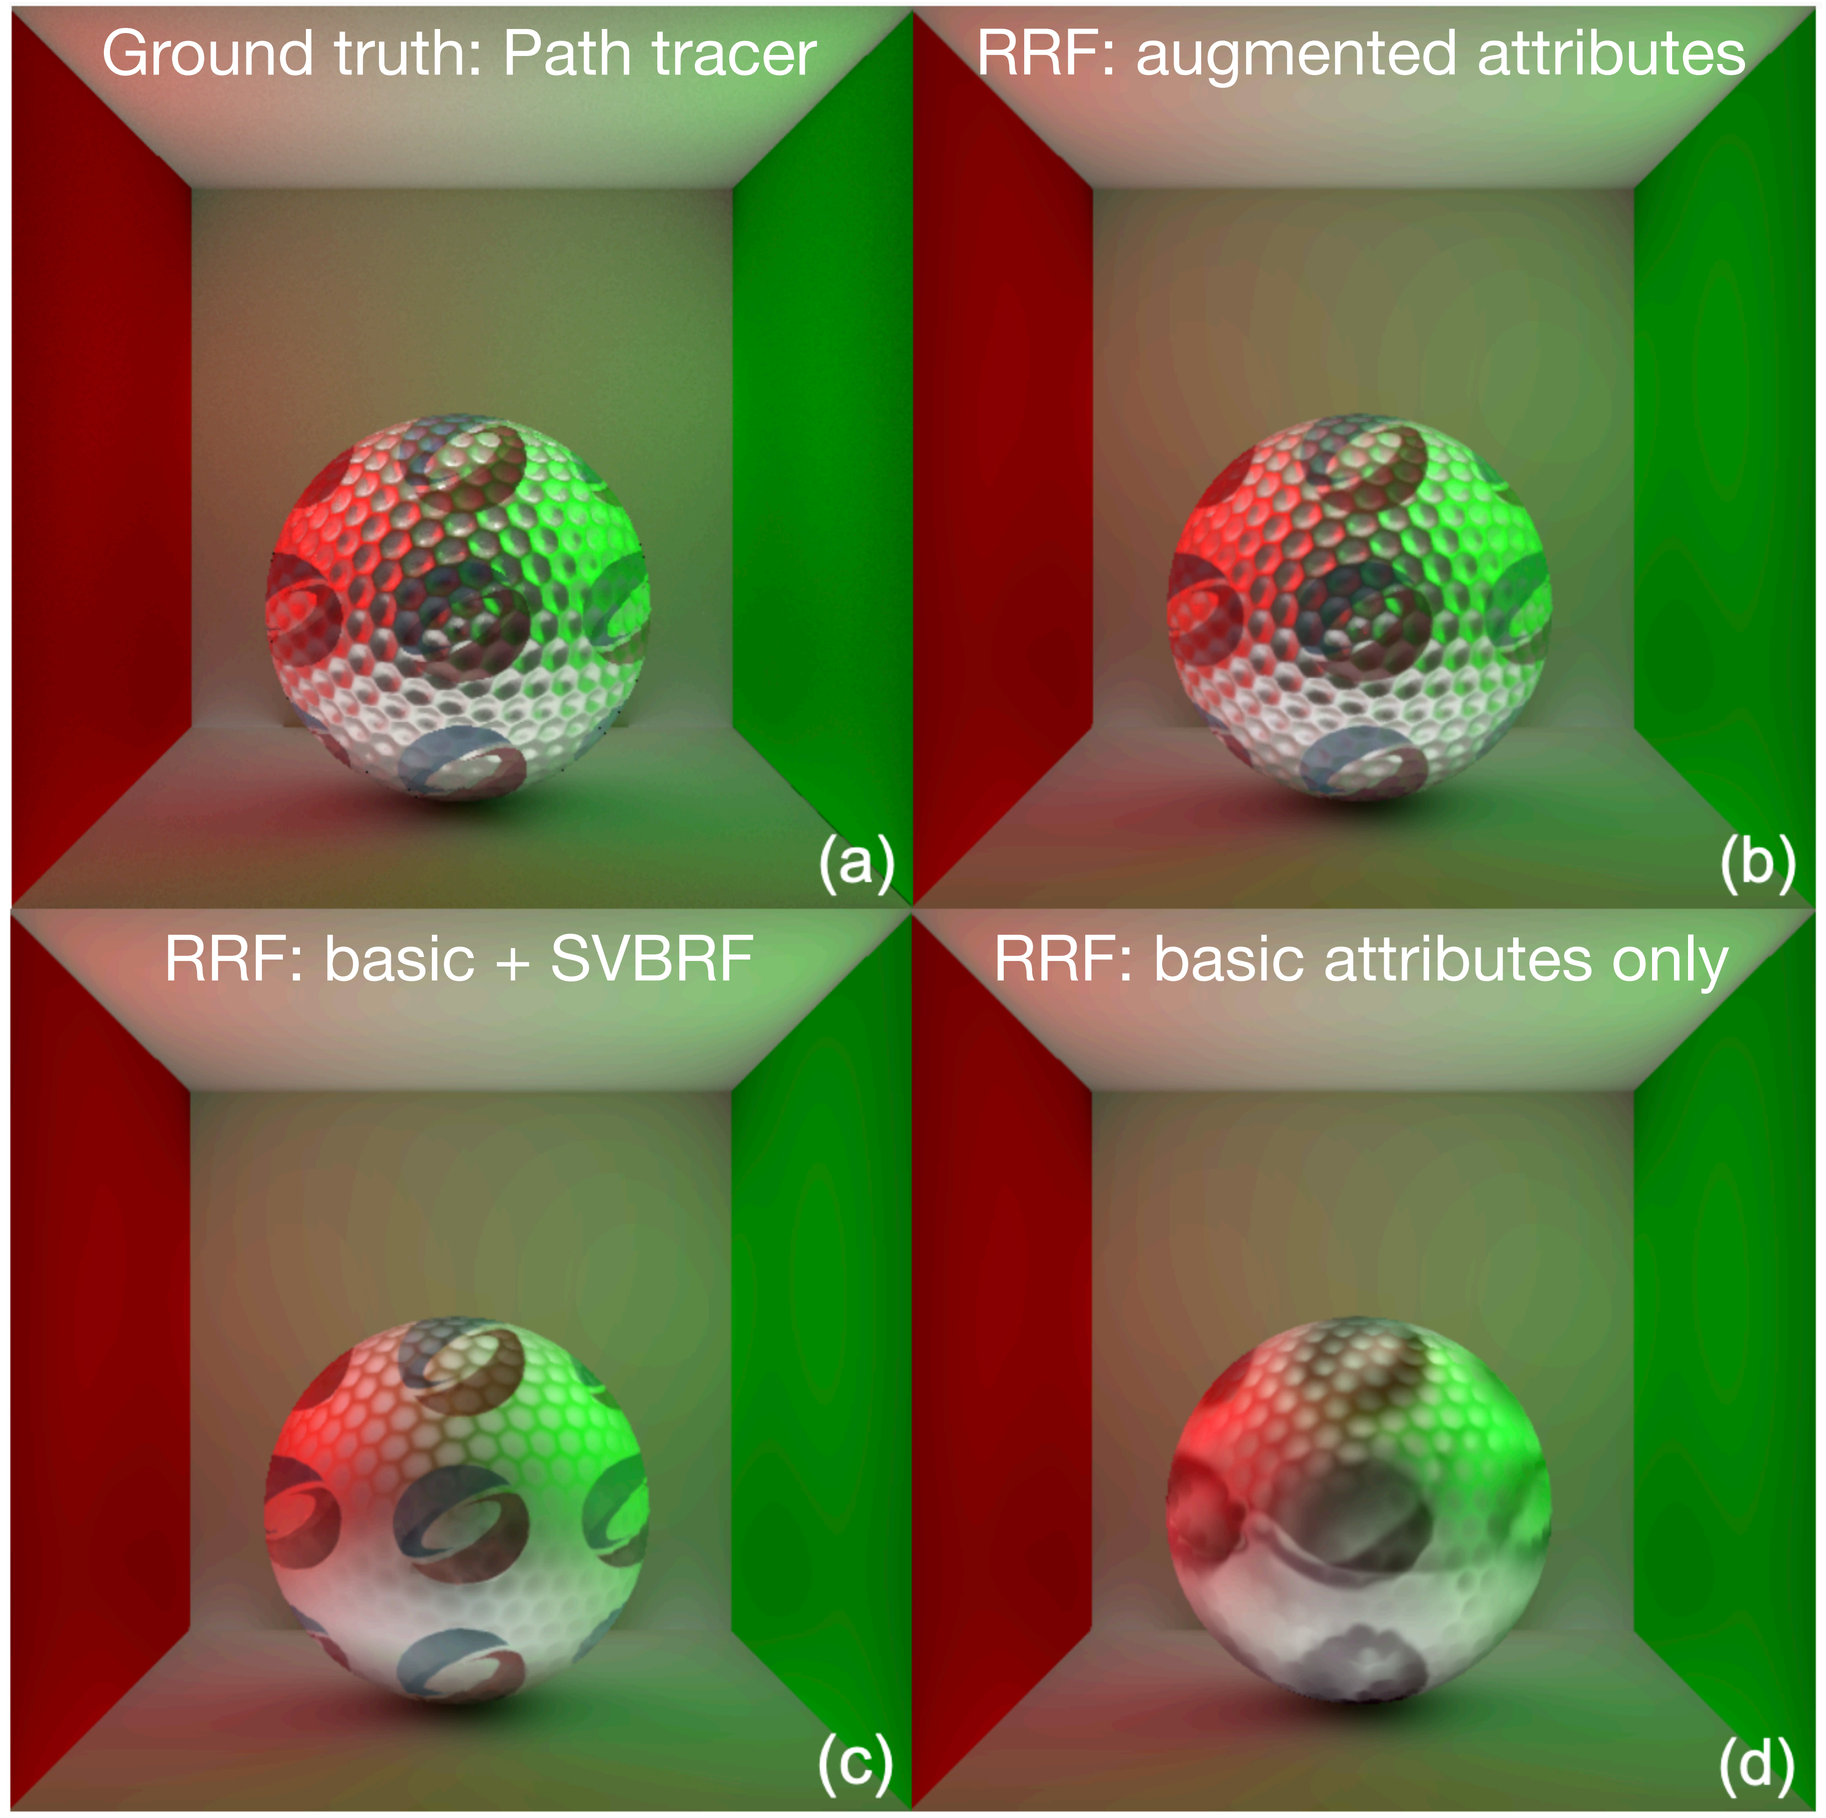

Only indirect components shown

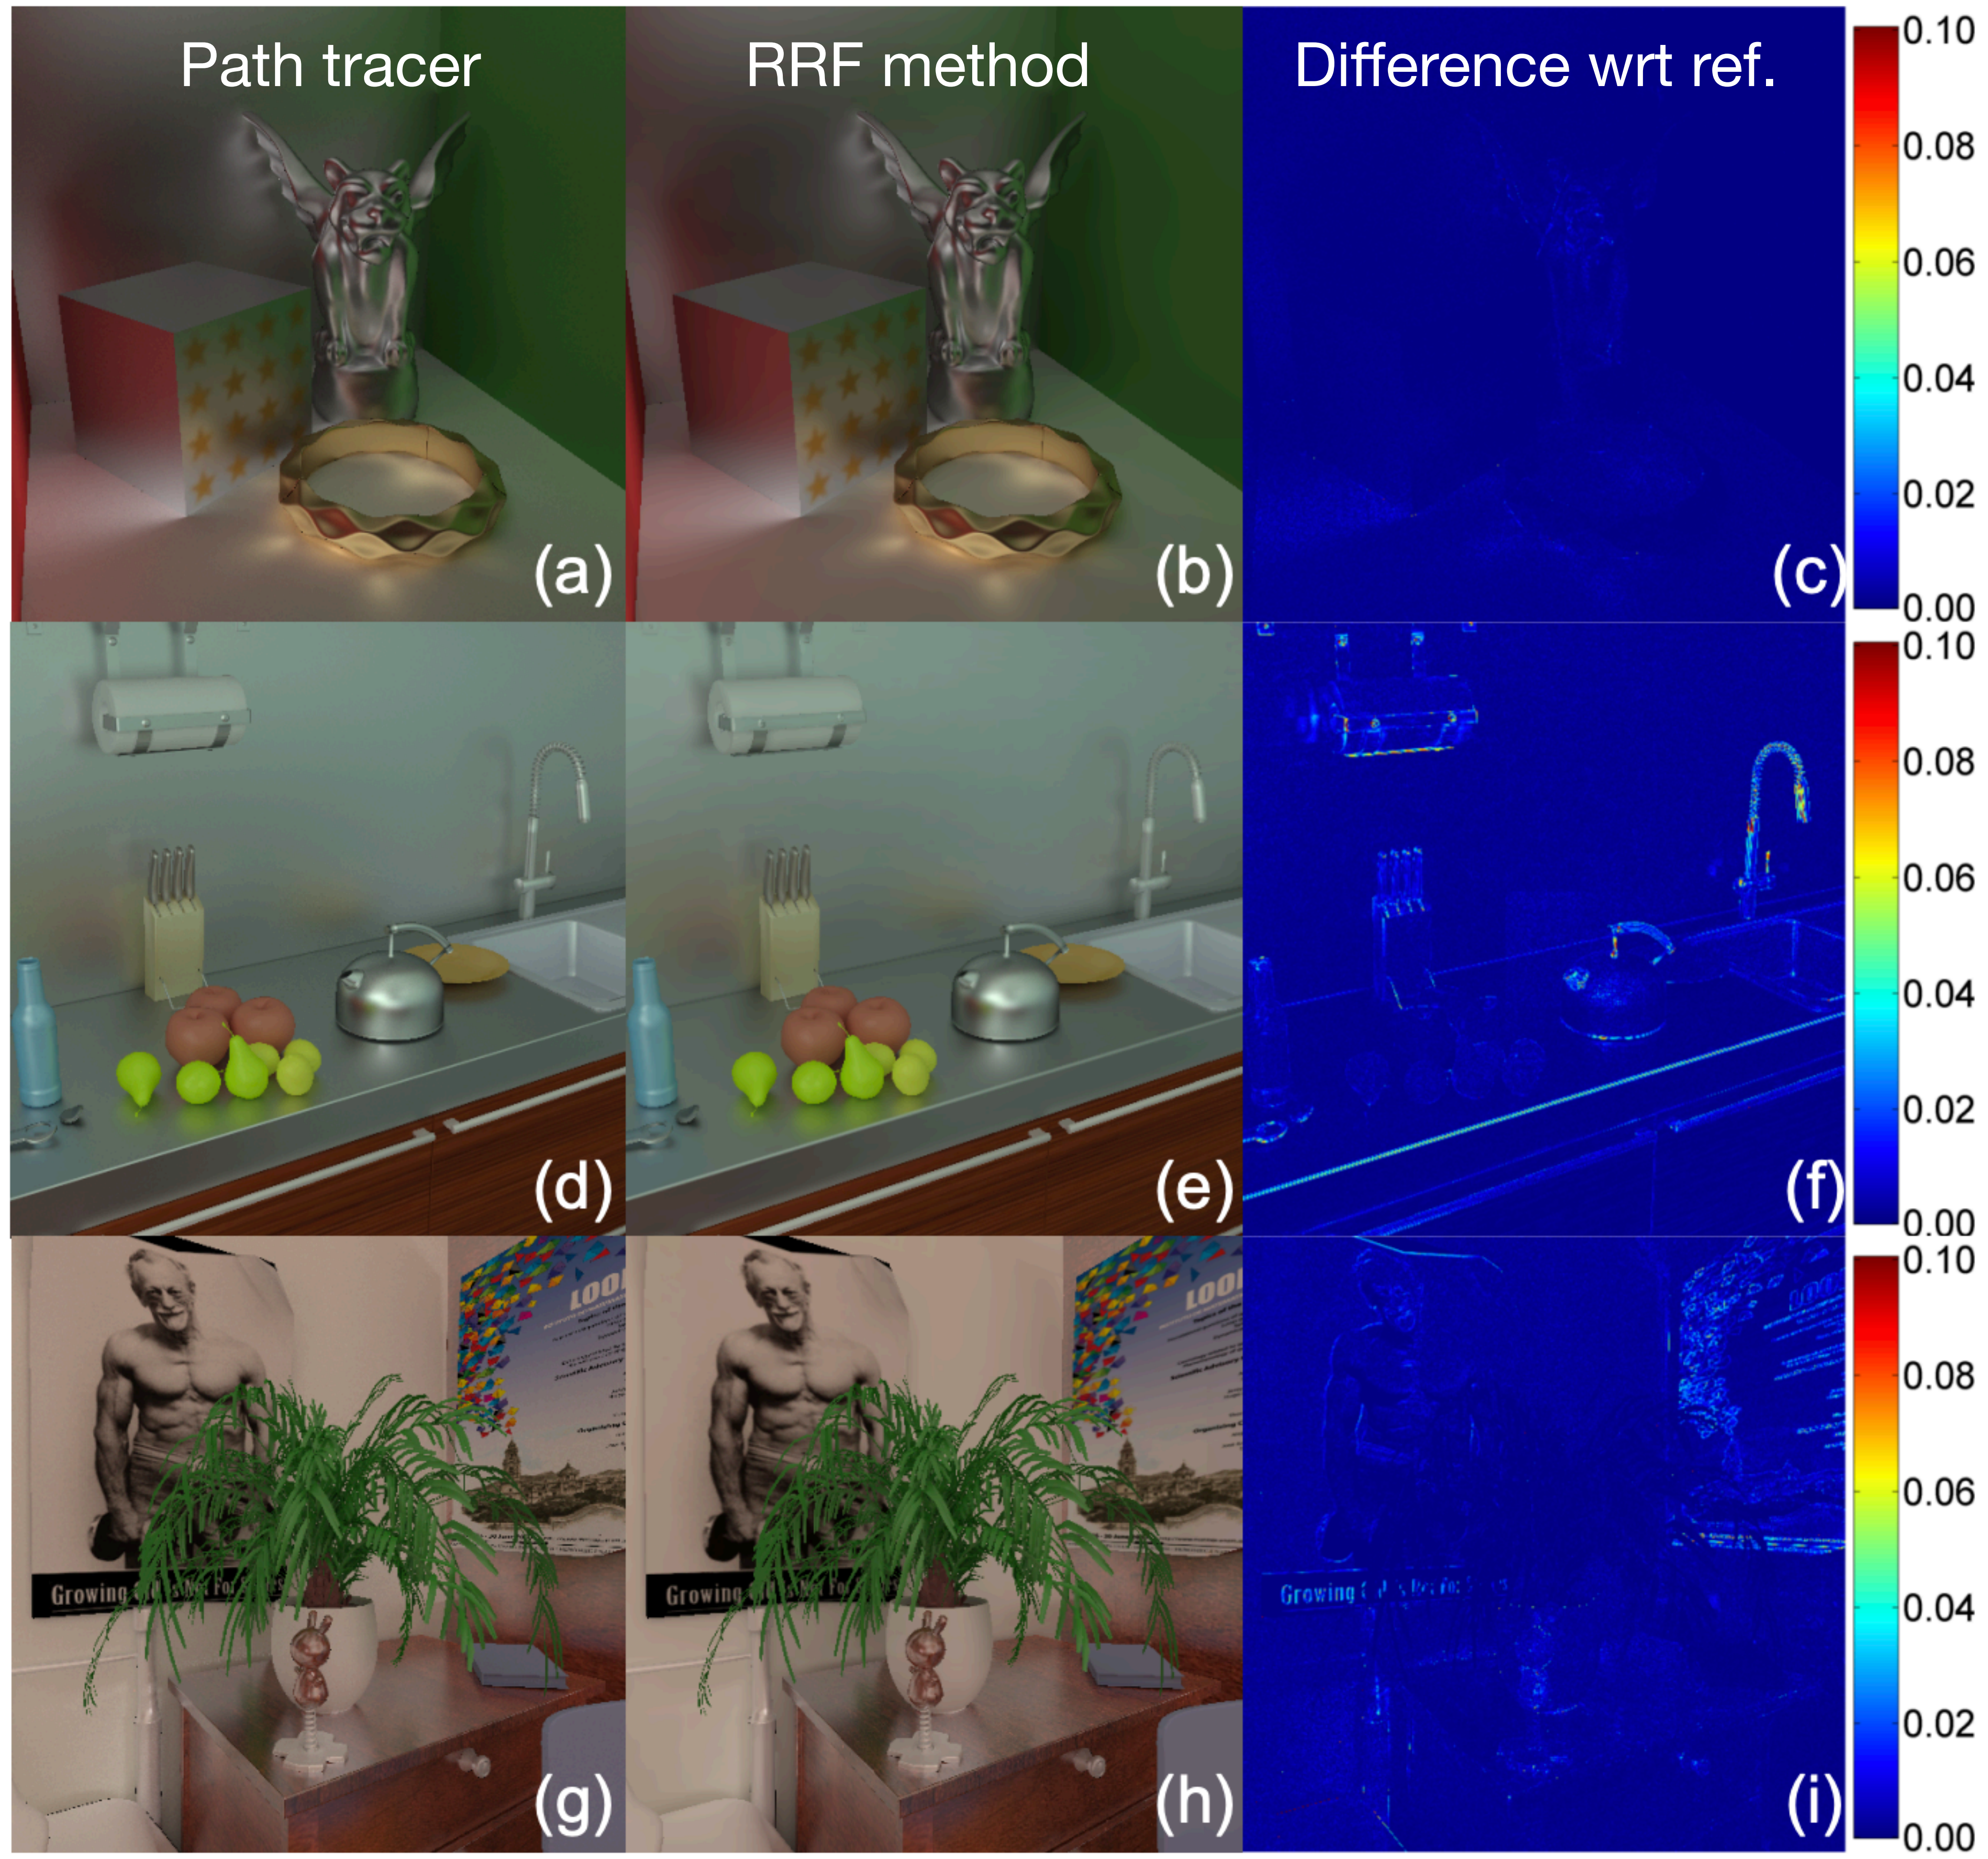

Contrastive learning path manifolds

Contrastive path learning
SIGGRAPH 2021

Converts reference pixel colors to dense pseudo labels for light paths

A convolutional path-embedding network:

- induces a low-dimensional manifold of paths by iteratively clustering intra-class embeddings,

- while discriminating inter-class embeddings using gradient descent

Previous reconstruction methods

SIGGRAPH 2021

Previous methods use auxiliary features at first bounce: normals, depth, texture

Some methods use indirect features (manually) for specular or non-specular

Challenge

A representation of light propagation is inherently high-dimensional

Learning meaningful patterns between high-dimensional paths and reference images is still challenging:

— due to the low correlation and high sparsity of path samples

Recent methods show that deep neural networks often struggle to explore the sparse space

Proposed framework

Aims to extract compact and useful embeddings of high-dimensional path features to remedy the sparsity of path space

Leverages contrastive approaches of existing deep architectures:

— which cluster input data foe tasks like classification and regression

Path space contrastive learning

Sample-based convolutional network
(transforms path descriptors
to a low-dimensional space)

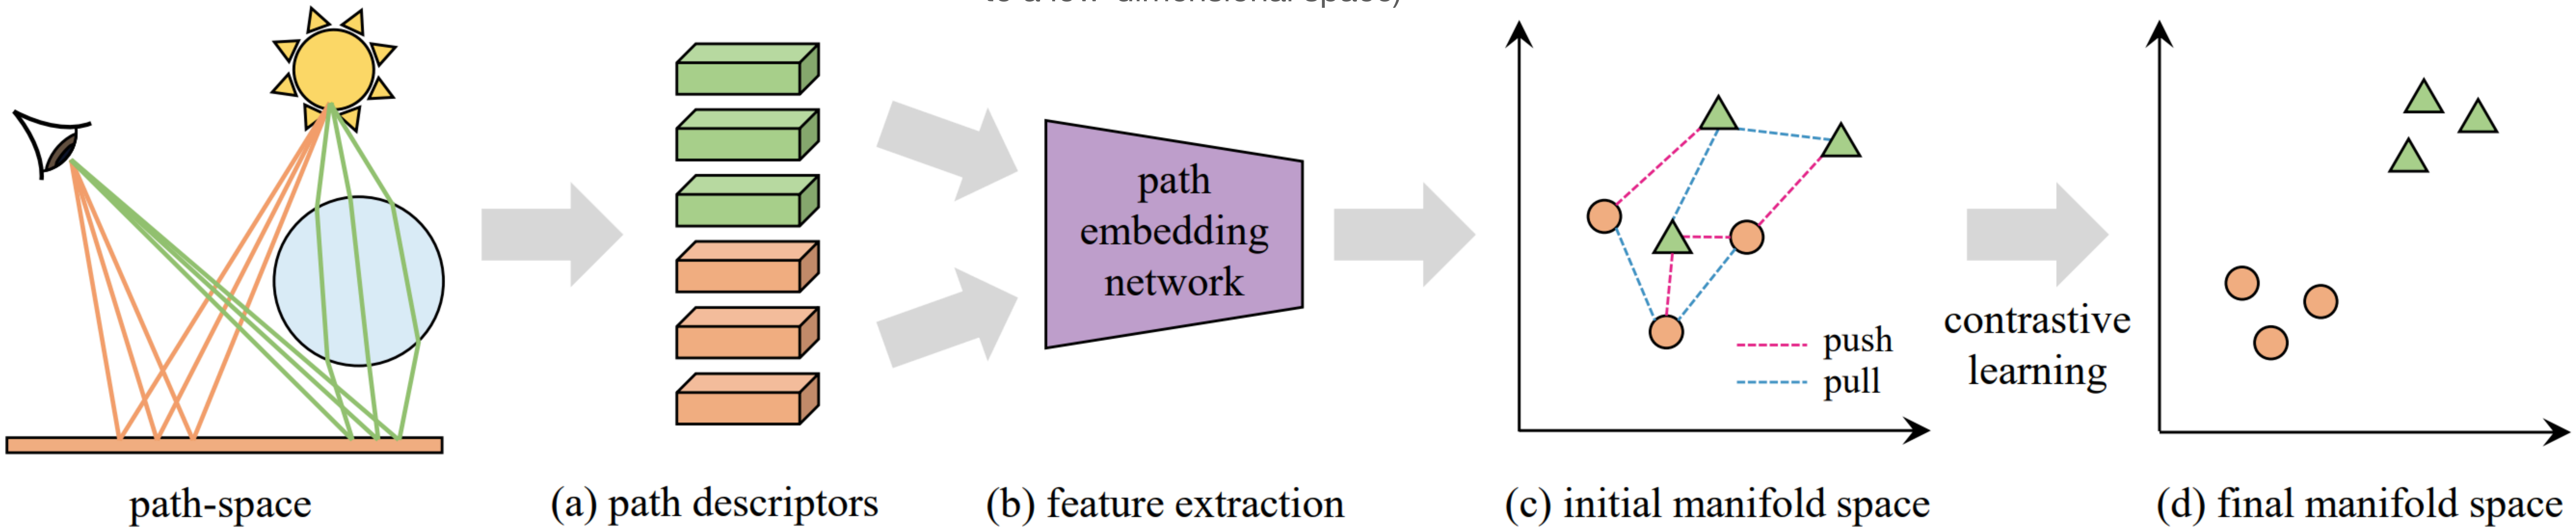

Extracts a path descriptor:
a sequence of the path's
radiometric quantities
at each vertex

Poorly-structured manifold space

Use reference colors to cluster
paths with similar pixel color

Manifold vs. Regression learning

This paper proposes direct sample-to-sample correlation to discriminate overlapped path distributions

Previous methods:

— image-space regression: learns correlation between input and target *pixels*

— sample-space model: learns correlation between input *samples* and target *pixels*

Path space contrastive learning

Joint manifold-regression training framework

Results: Path space contrastive learning

Neural radiance fields

NeRF

Mildenhall et al. 2020

NeRF pipeline

Mildenhall et al. 2020

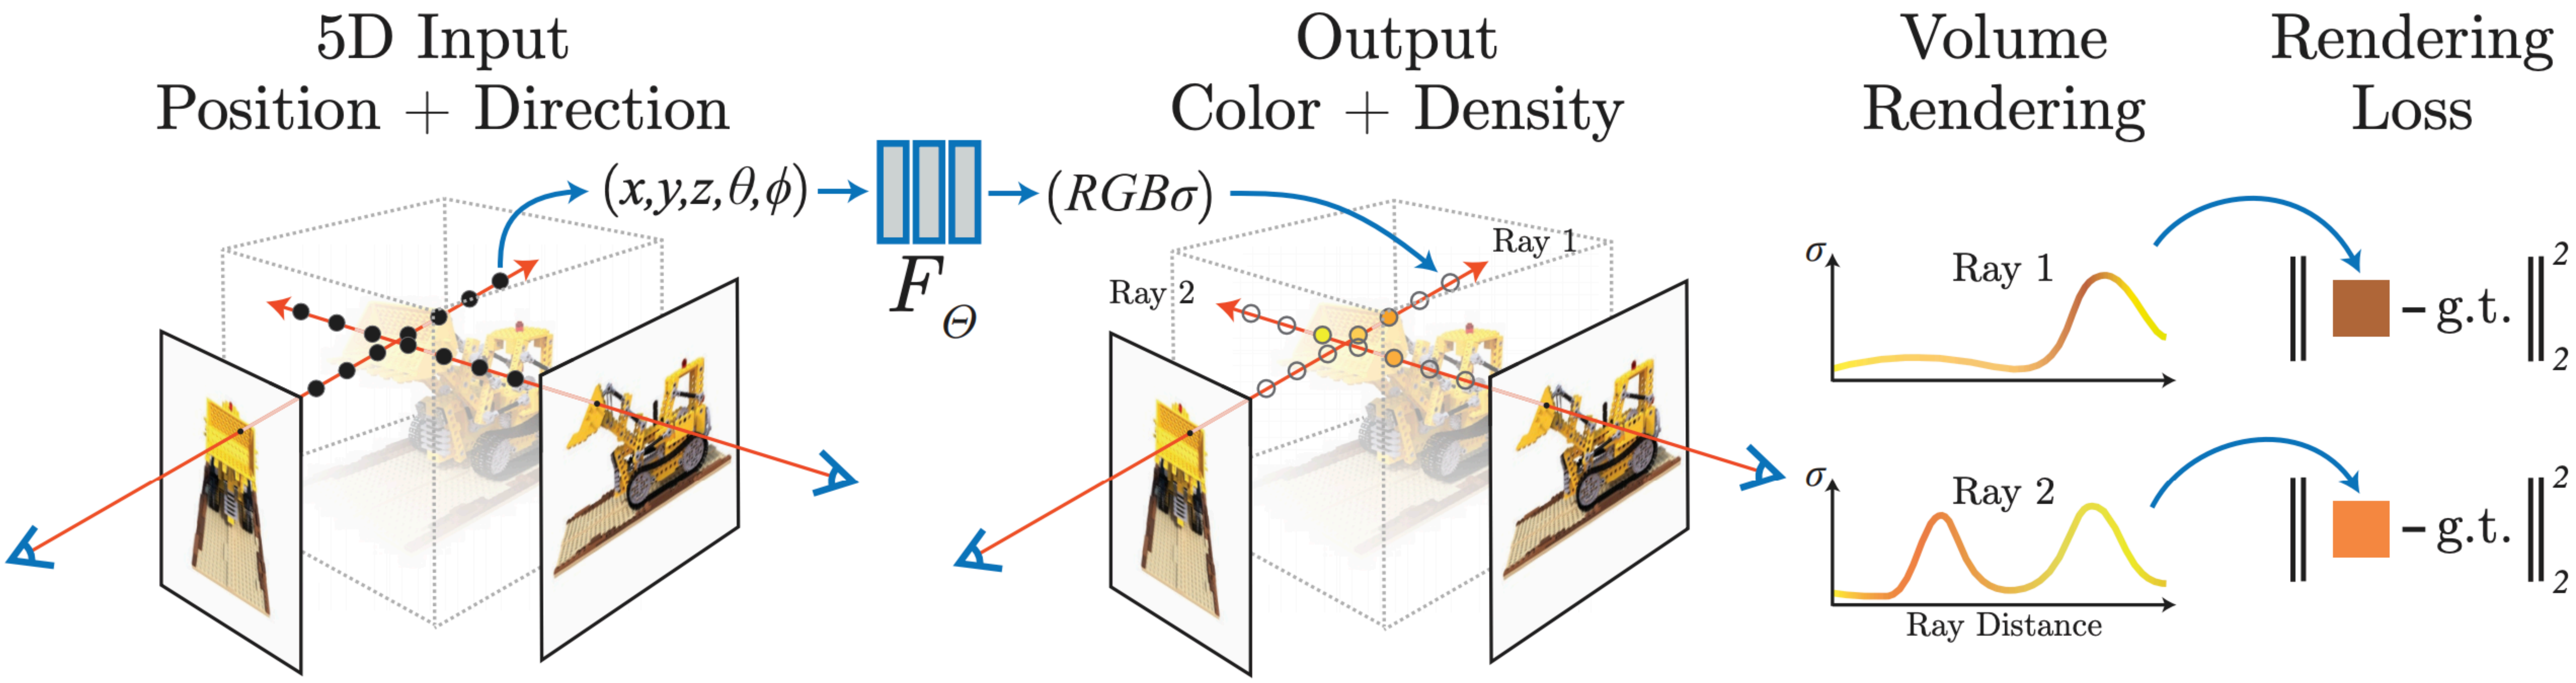

Mildenhall et al. 2020

Mildenhall et al. 2020

Neural control variates

What are control variates?

Recap

$$I = \int f(x)dx$$

What are control variates?

Recap

$$I = \int f(x) dx$$

$$I = \int \left(f(x) + g(x) - g(x) \right) dx$$

What are control variates?

Recap

$$I = \int f(x)dx$$

$$I = \int \left(f(x) + g(x) - g(x)\right) dx$$

$$I = \int g(x)dx + \int \left(f(x) - g(x)\right) dx$$

What are control variates?

Recap

$$I = \int g(x)dx + \int \left(f(x) - g(x)\right) dx$$

$g(x)$: control variate

What are control variates?

Recap

$$I = \int g(x)dx + \int \left(f(x) - g(x)\right) dx$$

$g(x)$: control variate

$$I = G + \int \left(f(x) - g(x)\right) dx$$

What are control variates?

Recap

$$I = \int g(x)dx + \int \left(f(x) - g(x)\right) dx$$

$g(x)$: control variate

$$I = G + \int \left(f(x) - g(x)\right) dx$$

$$G = \int g(x)dx$$

What are control variates?

Recap

$$I = \int g(x)dx + \int \left(f(x) - g(x)\right) dx$$

$g(x)$: control variate

$$I = G + \int \left(f(x) - g(x)\right) dx$$

$$G = \int g(x)dx$$

$$I = G + \sum_{k=1}^{N} \left(f(x_k) - g(x_k)\right)$$

What are control variates?

Recap

$$I = \int g(x)dx + \int (f(x) - g(x))\, dx$$

$g(x)$: control variate

$$I = G + \int (f(x) - g(x))\, dx$$

$$G = \int g(x)dx$$

$$I = G + \sum_{k=1}^{N} (f(x_k) - g(x_k))$$

$$I = \alpha G + \sum_{k=1}^{N} (f(x_k) - \alpha g(x_k))$$

What are control variates?

How they impact the estimator?

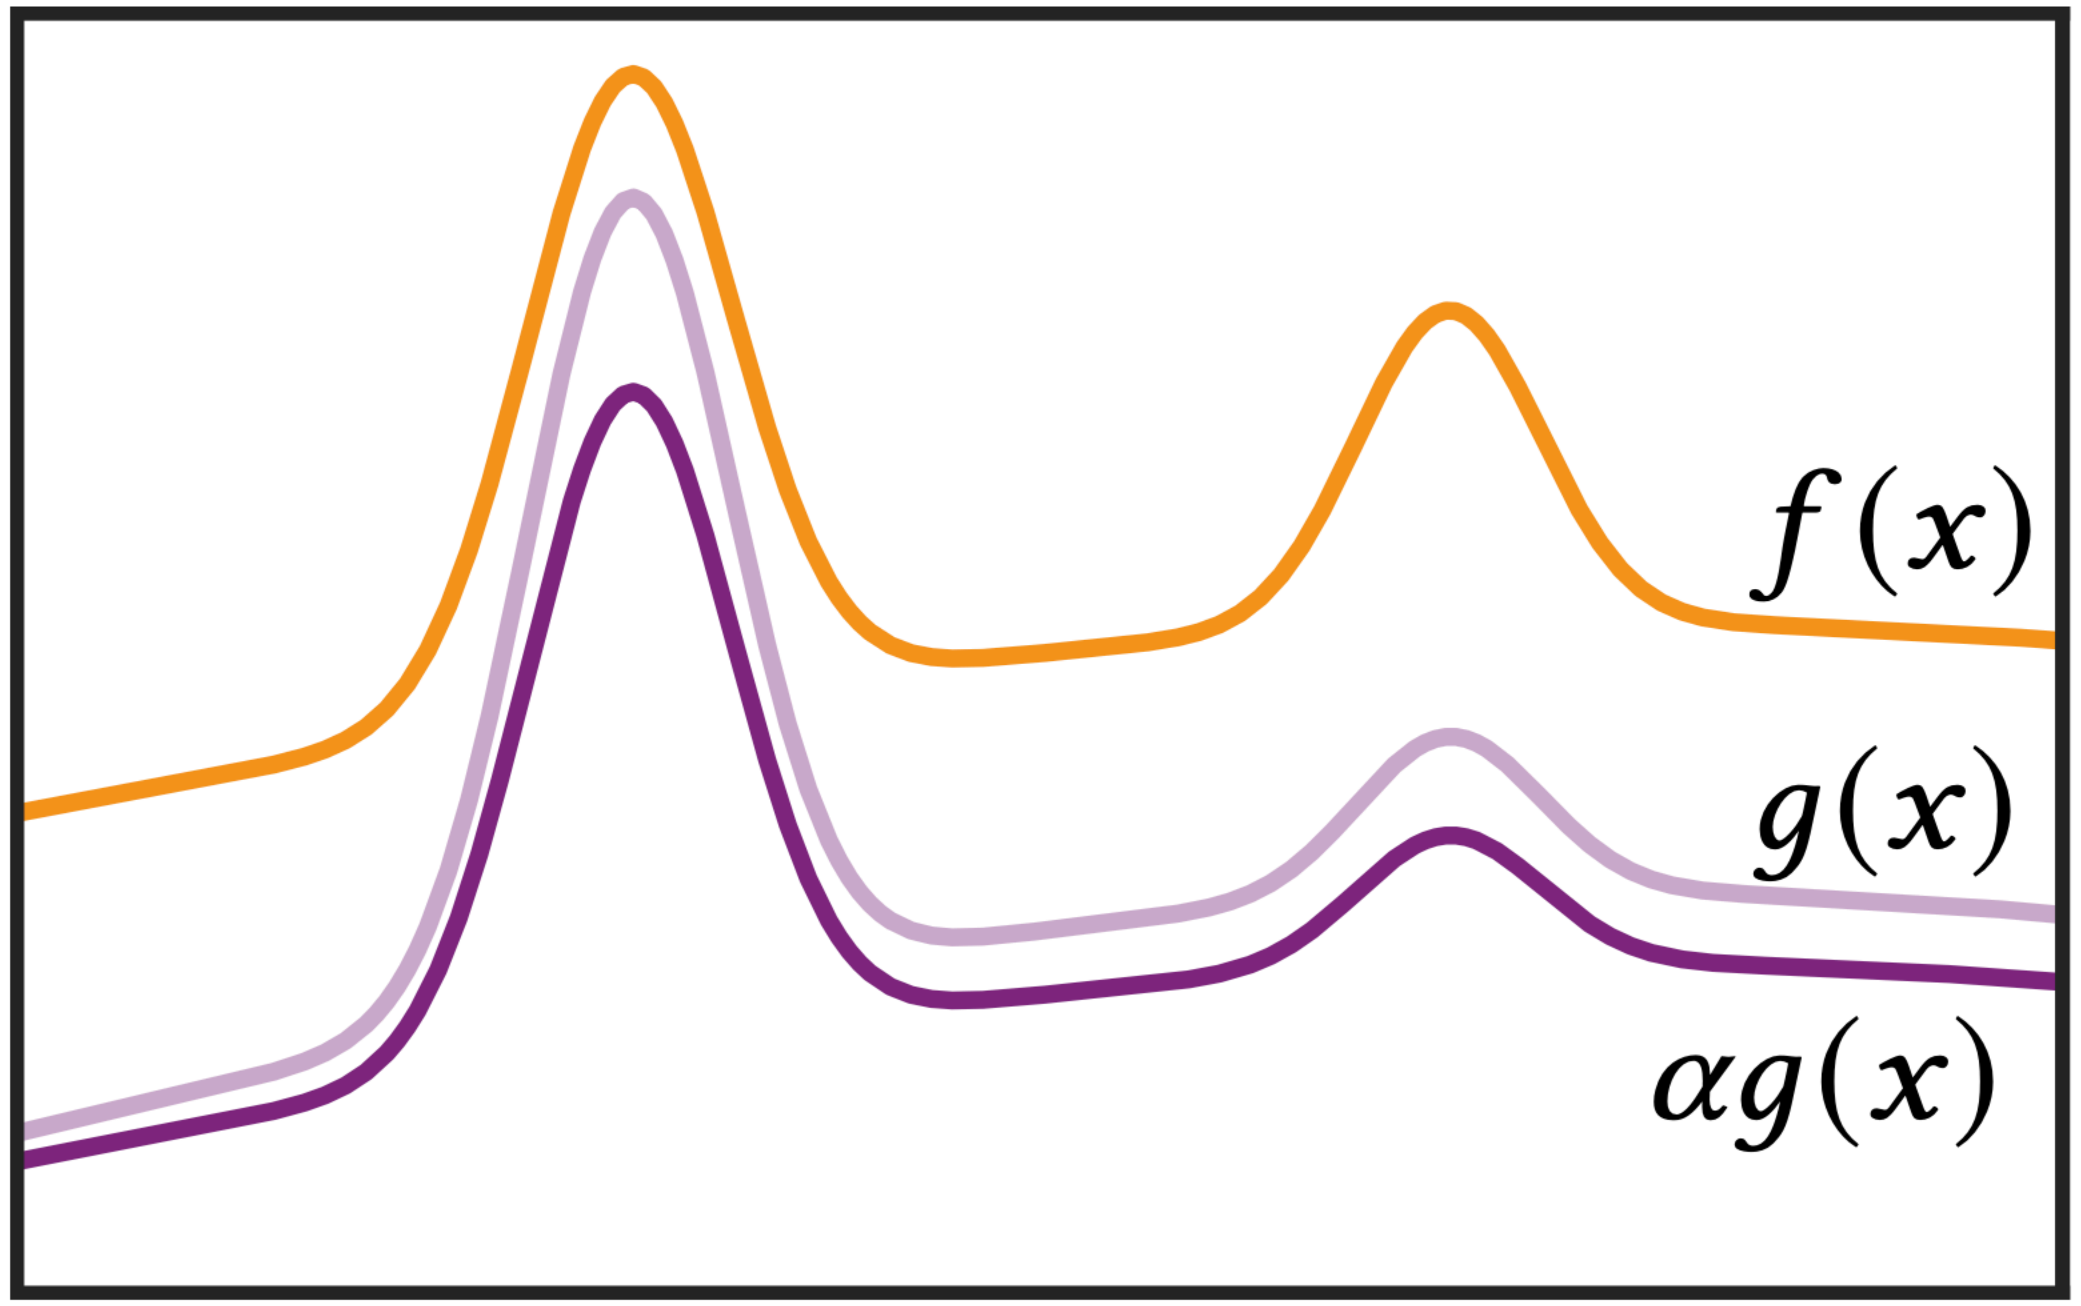

What are control variates?

How they impact the estimator?

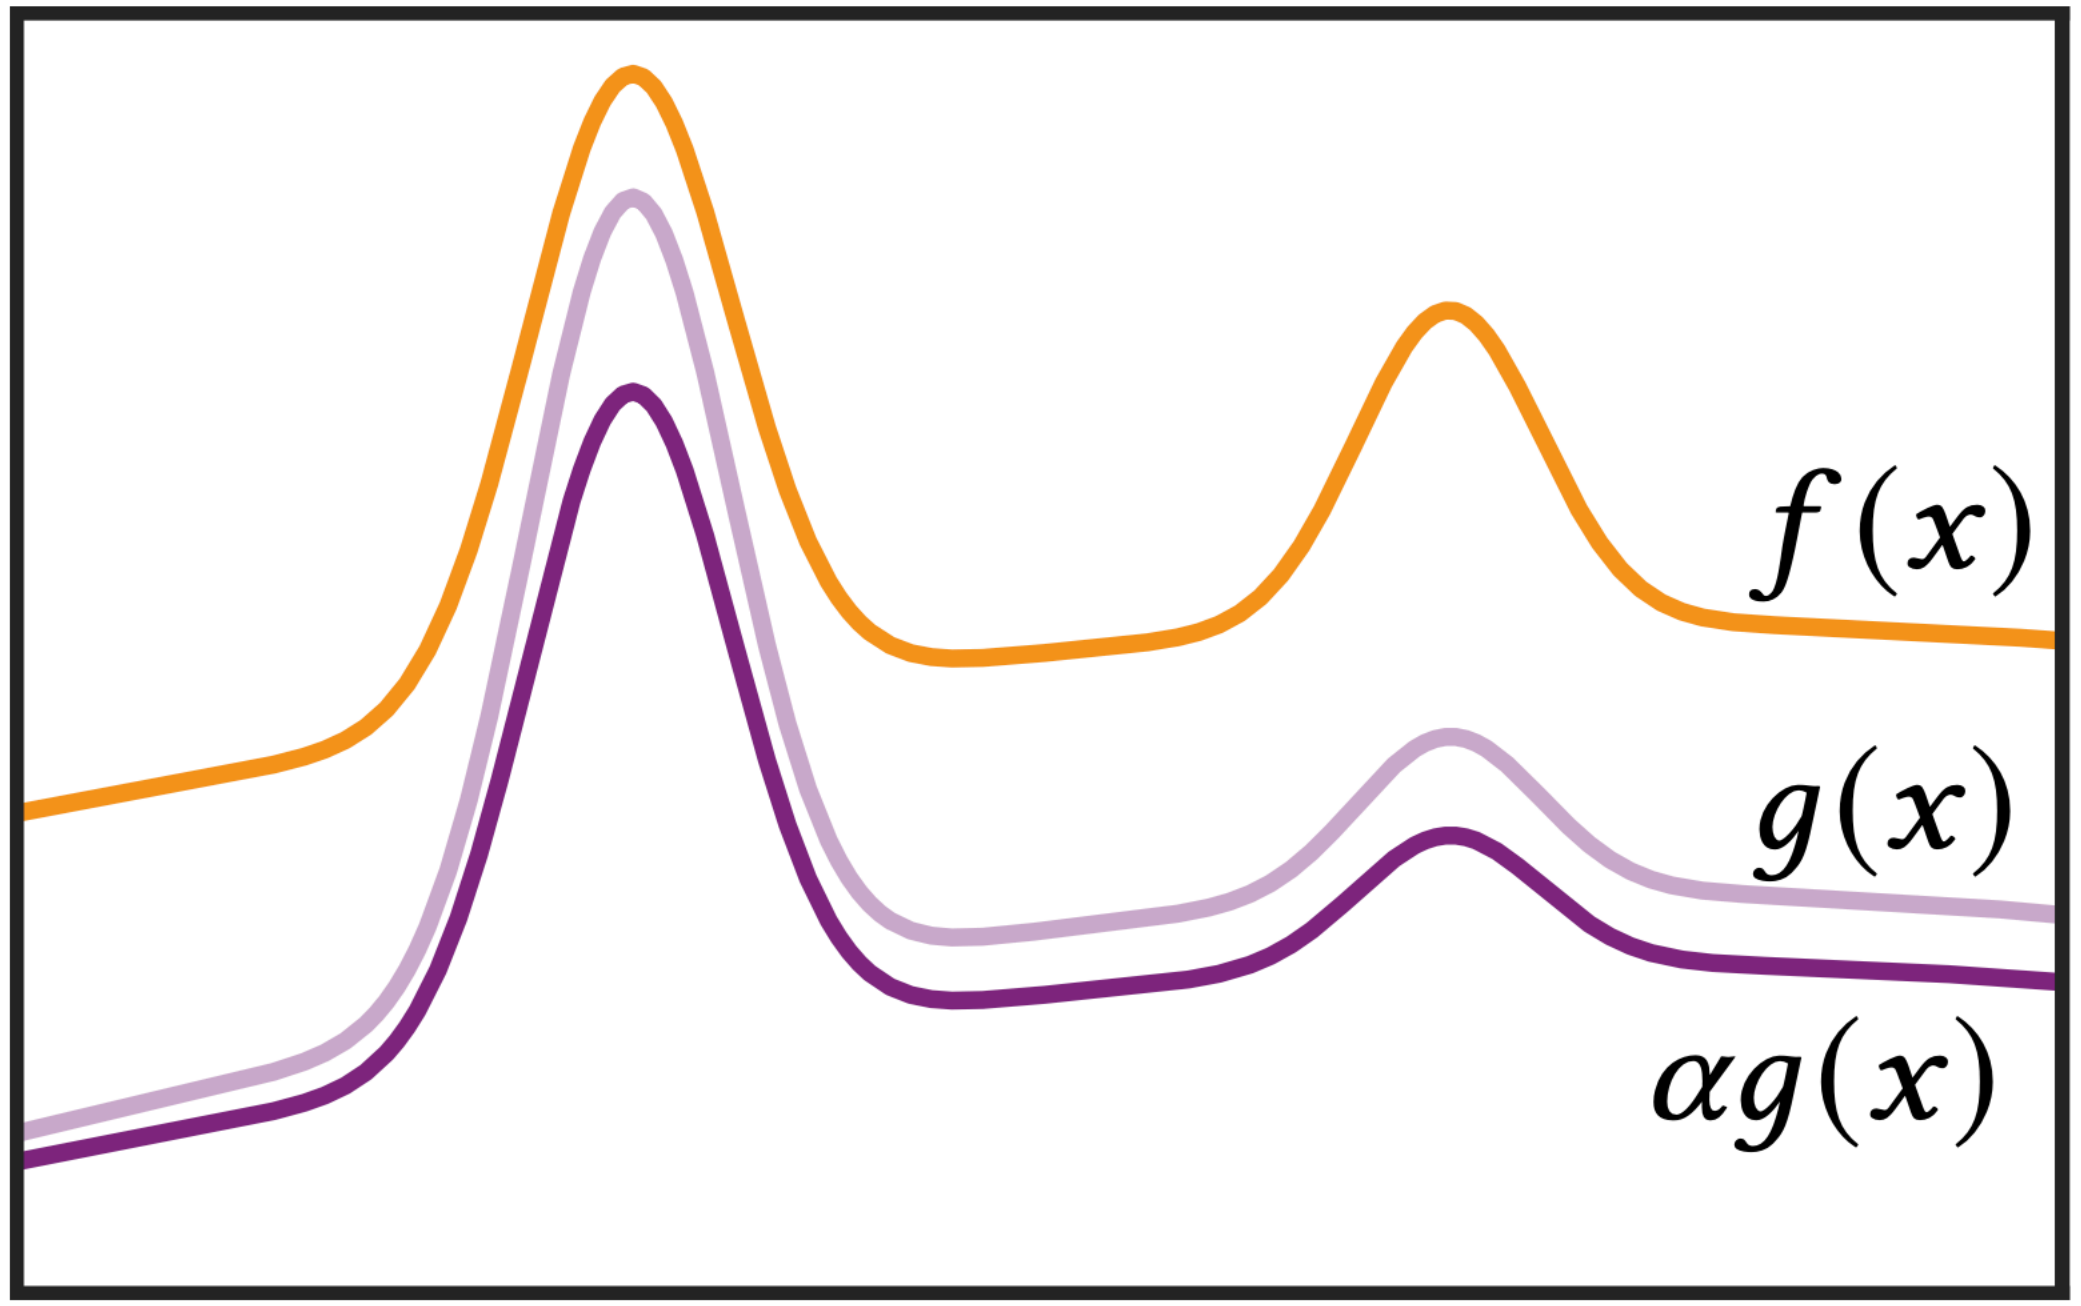

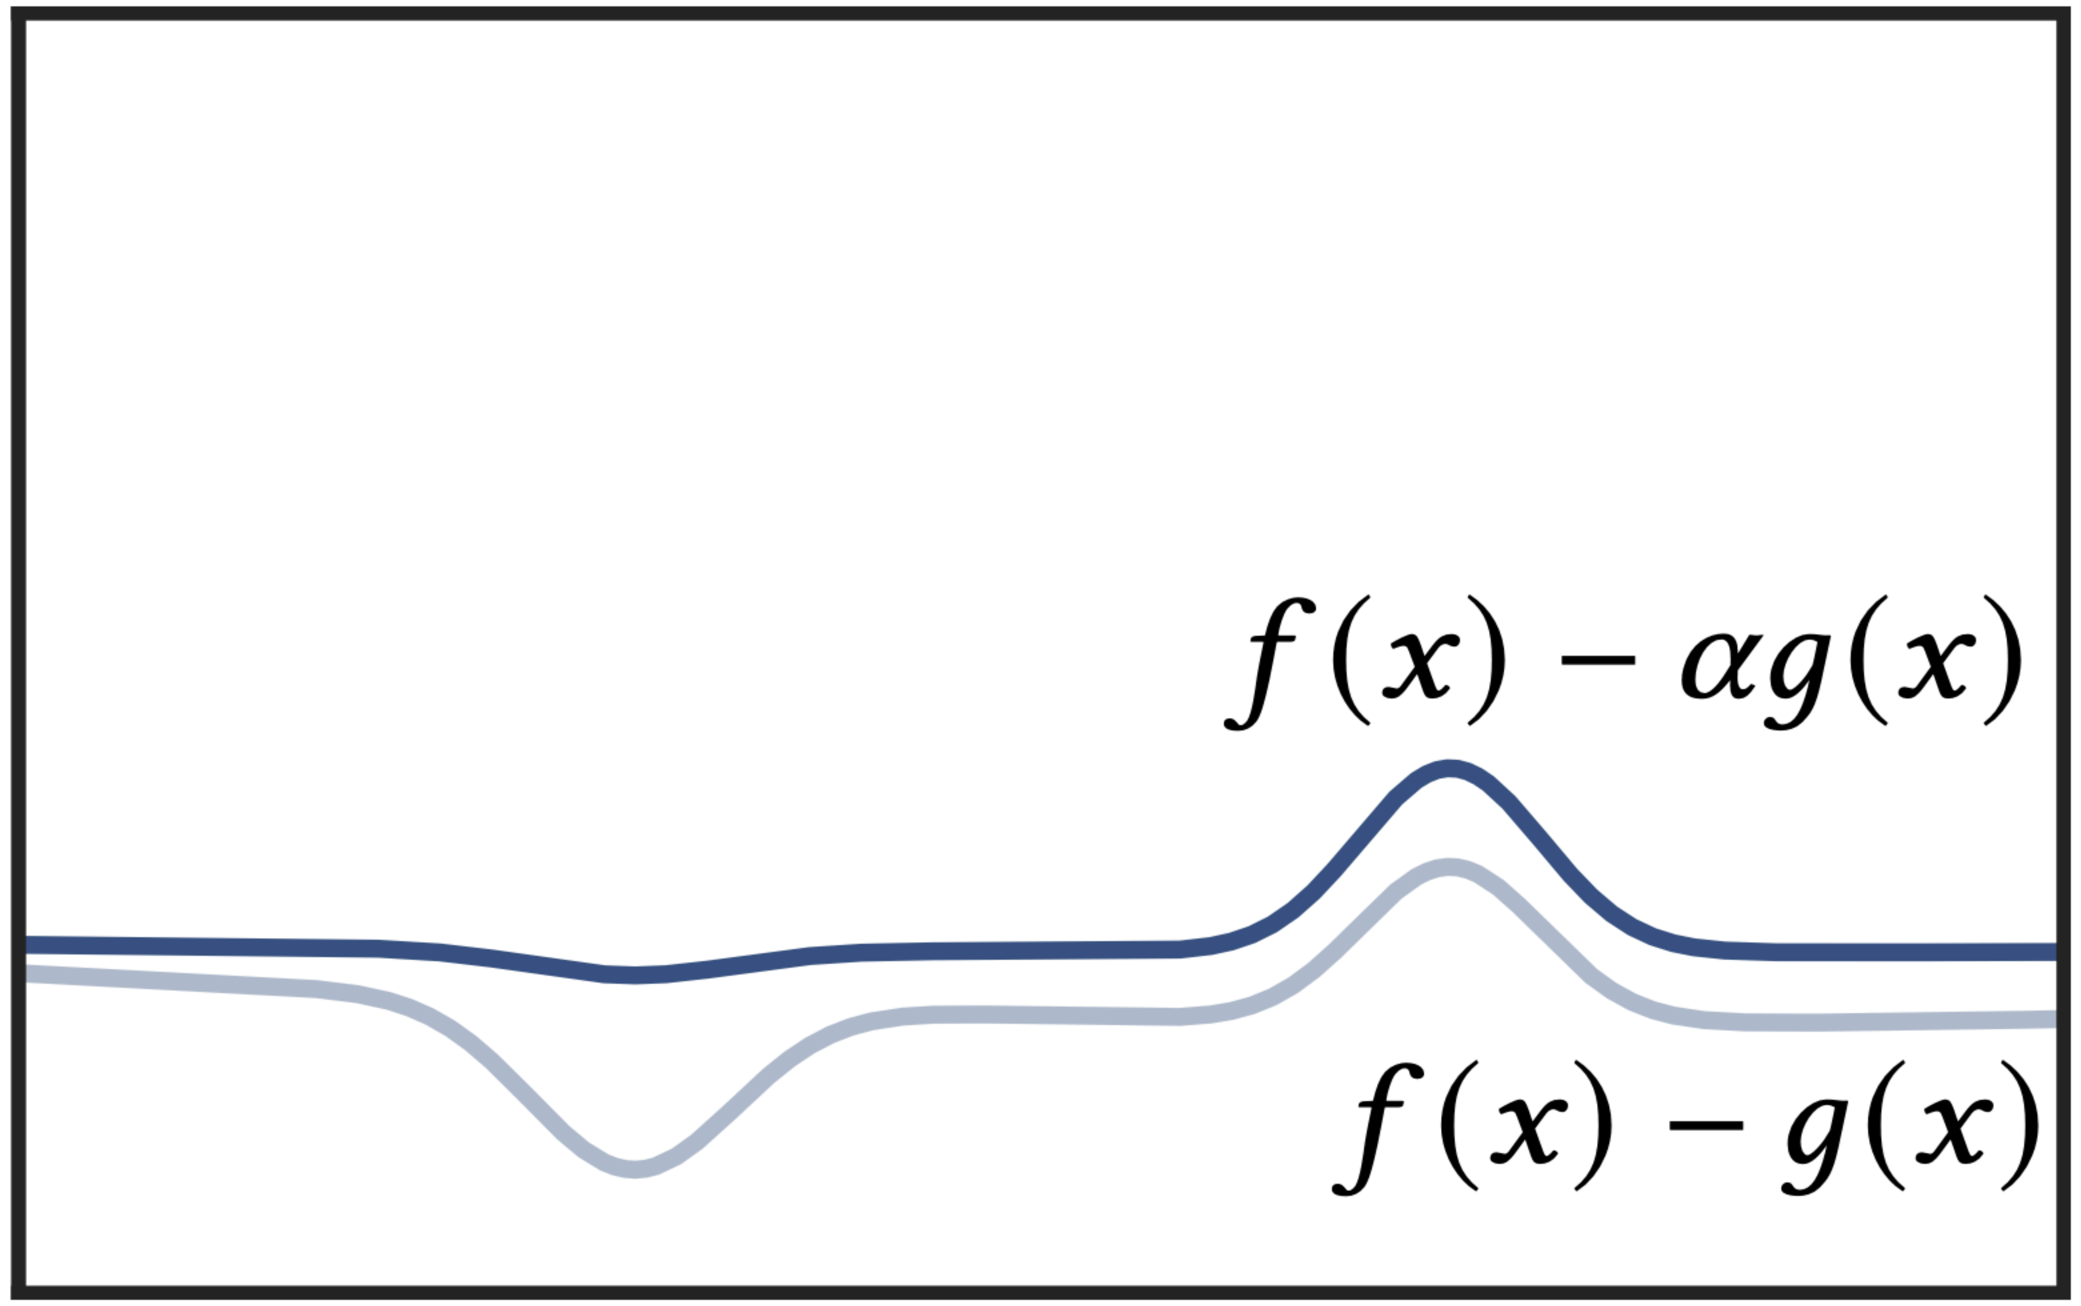


Render time: sometimes >100 cpu-hours

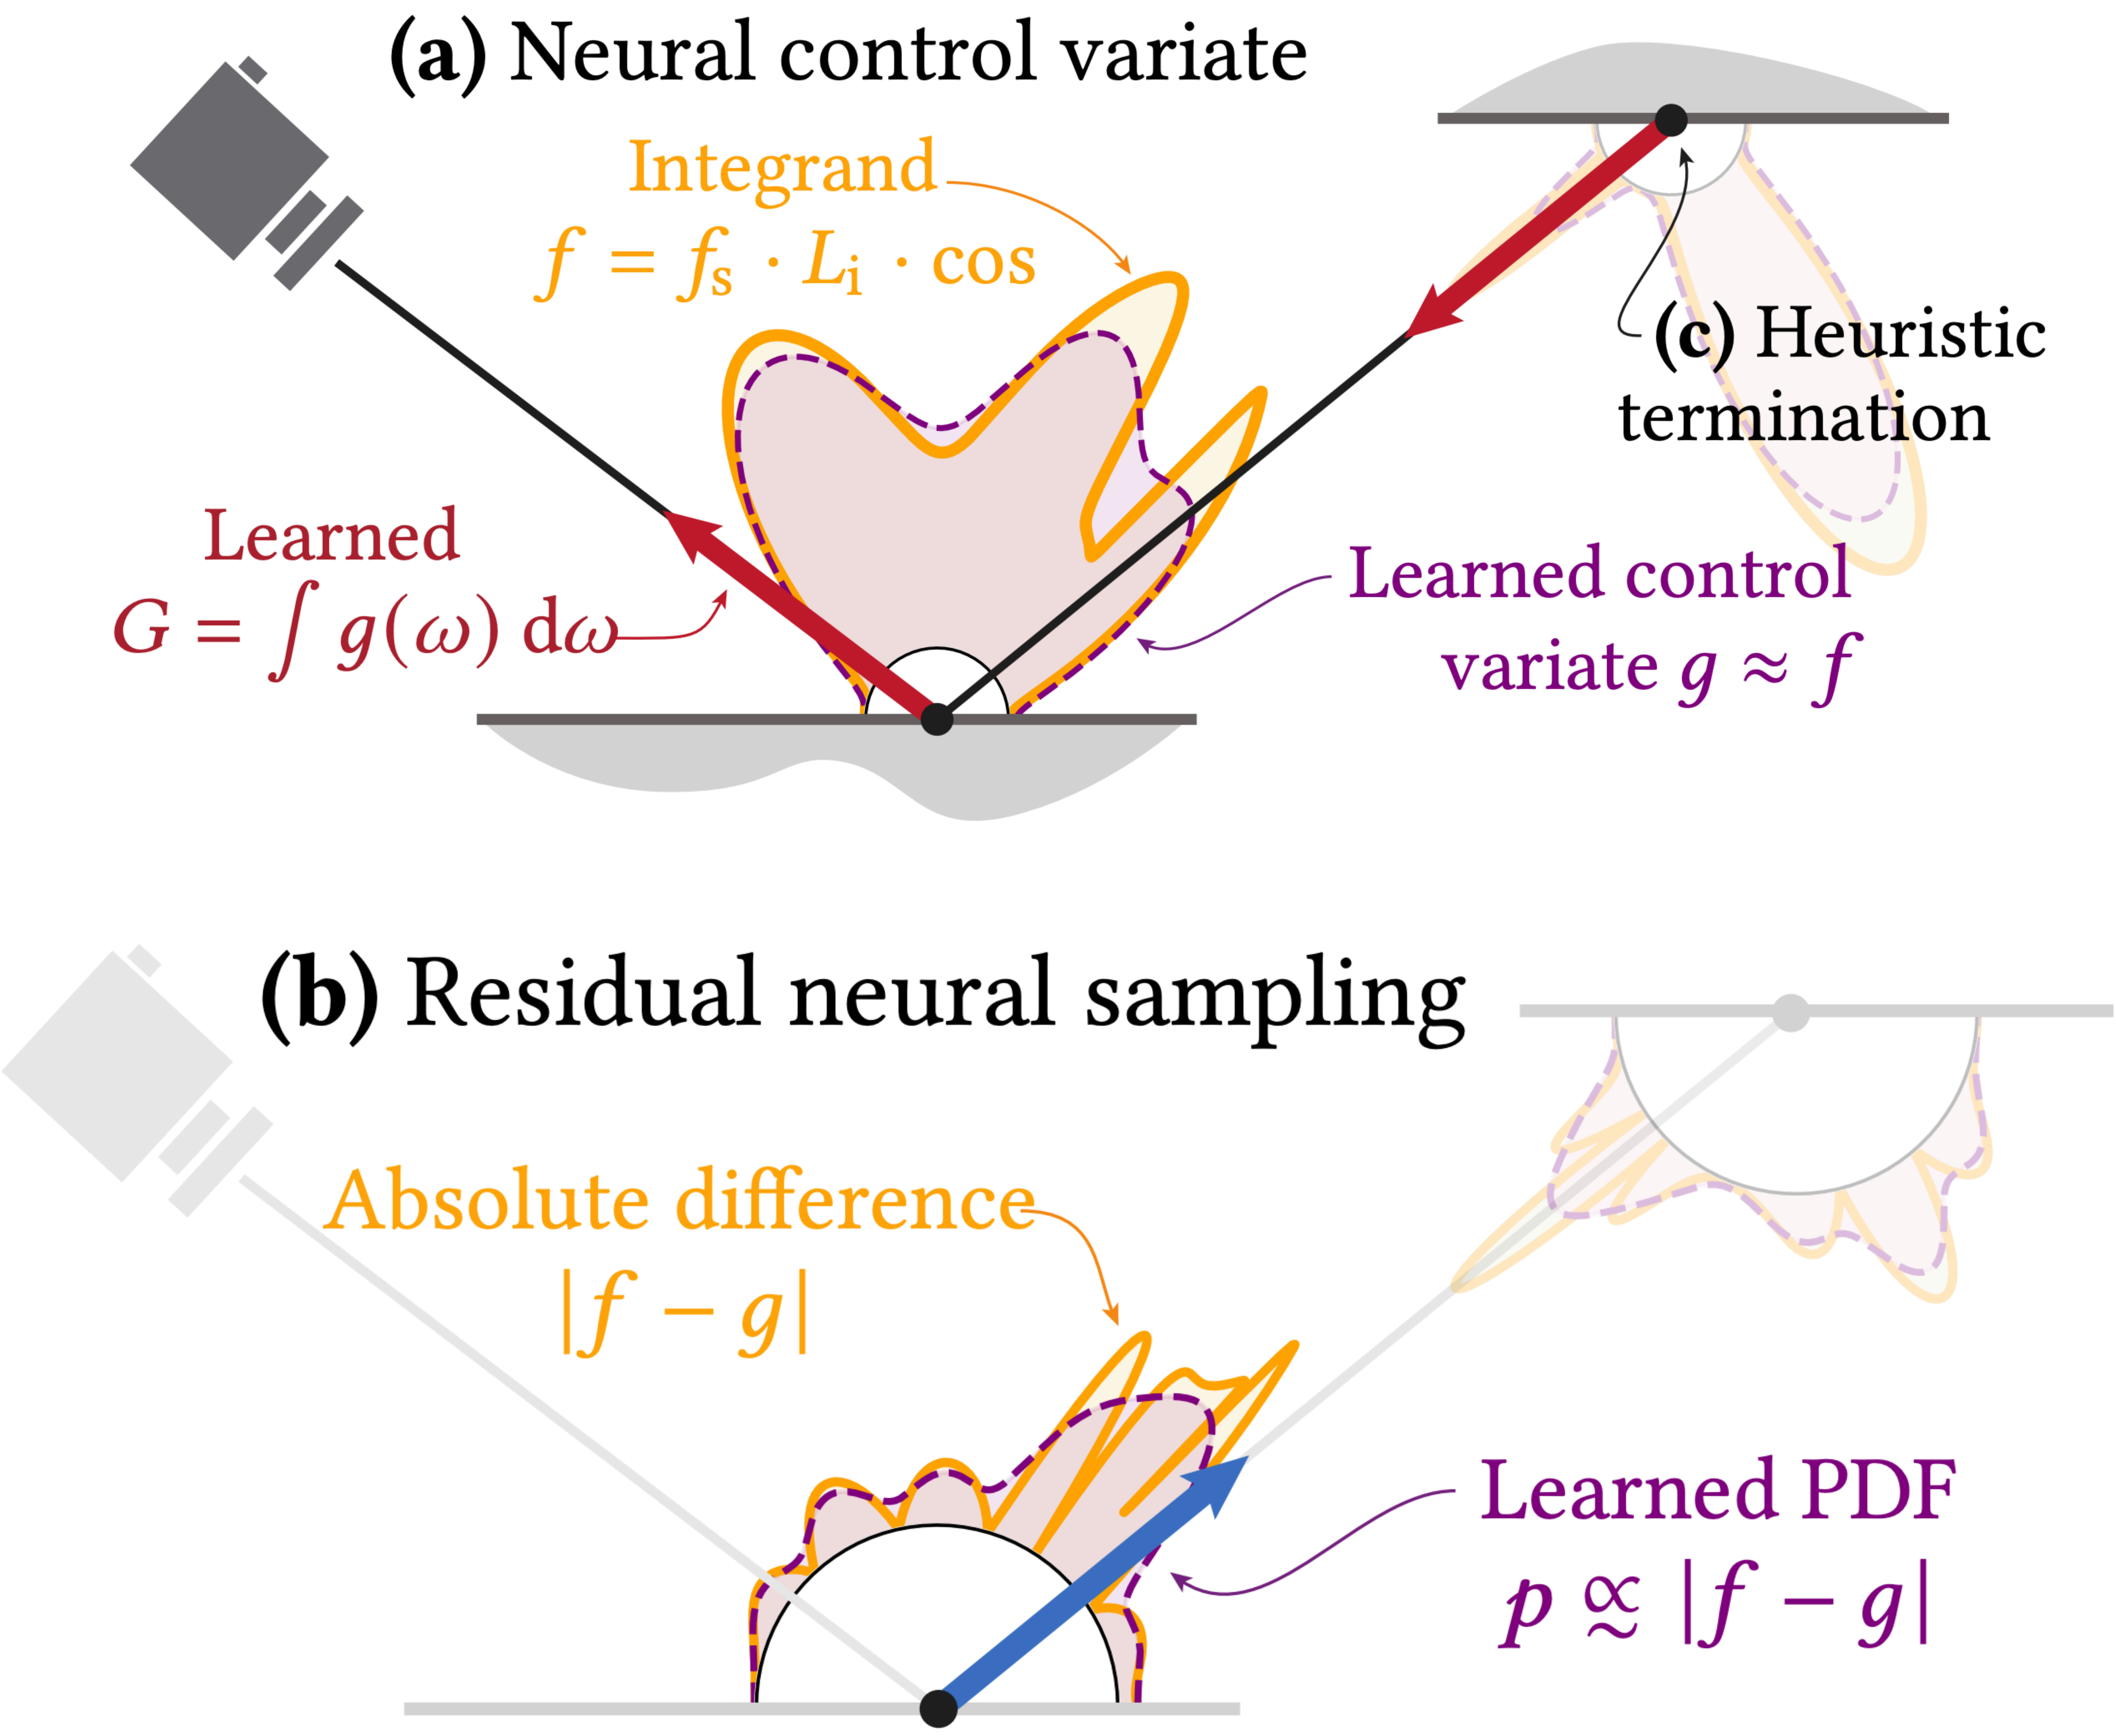
(a) Neural control variate
Integrand
$f = f_s \cdot L_i \cdot \cos$
Learned
$G = \int g(\omega)\,d\omega$
Learned control variate $g \approx f$
(c) Heuristic termination
(b) Residual neural sampling
Absolute difference
$|f - g|$
Learned PDF
$p \propto |f - g|$


Unbiased
Biased
PT
PPG
NIS++
NCV
NCV + heuristic
CV Integral
Reference
Bathroom
Veach Door

Spatially-correlated transmittance

Spatially-correlated media

Impact on rendering

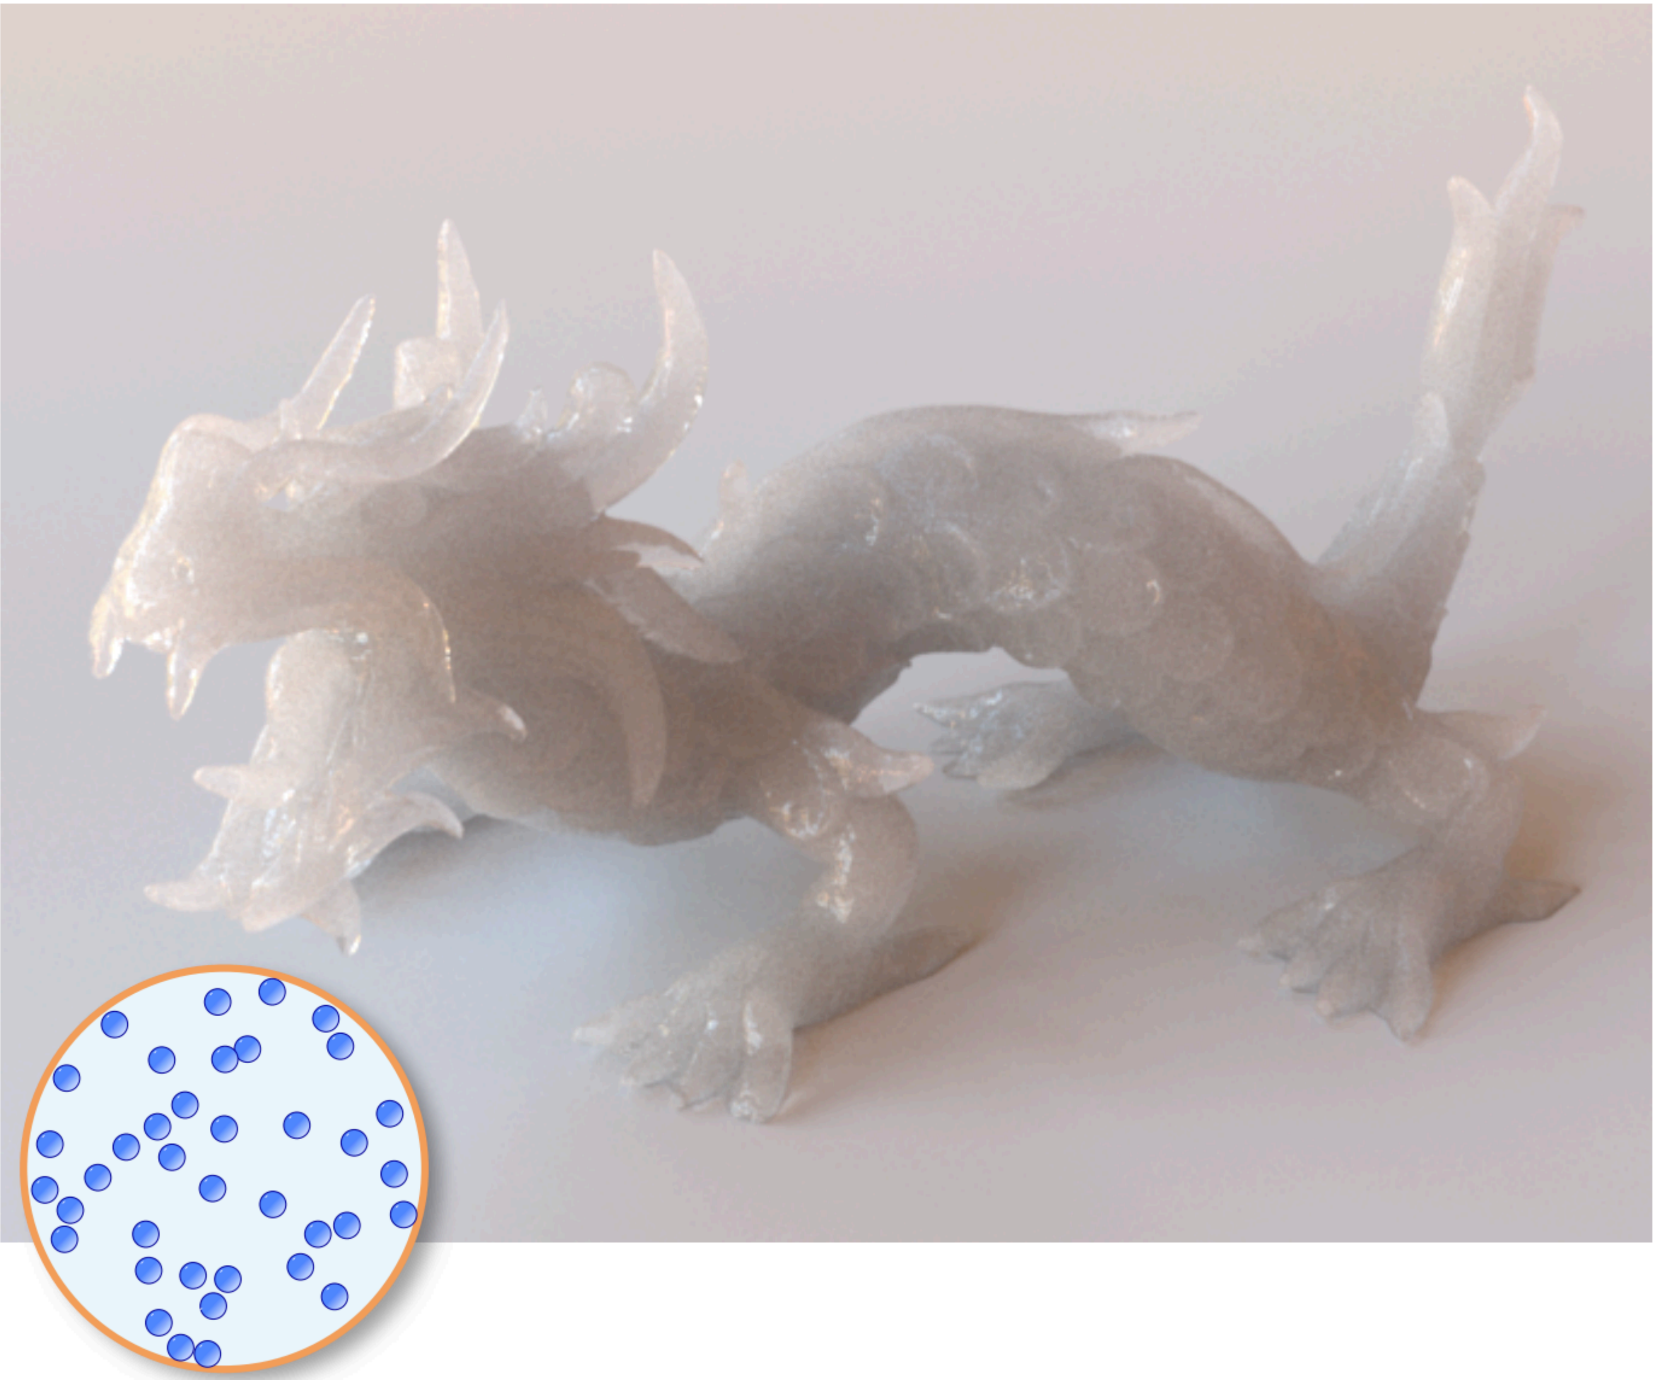

Uncorrelated media

Positively-correlated media

negatively-correlated media

Non-exponential transmittance

Transmittance function as a function of distance

Volumetric scene representations

Transmittance function as a function of distance

Volumetric representations are useful for complex scenes

Becoming popular for level of detail and scene reconstruction

Traditional exponential transmittance model cannot capture correlations in visibility across volume elements

Volumetric scene representations

Transmittance due to different visibility

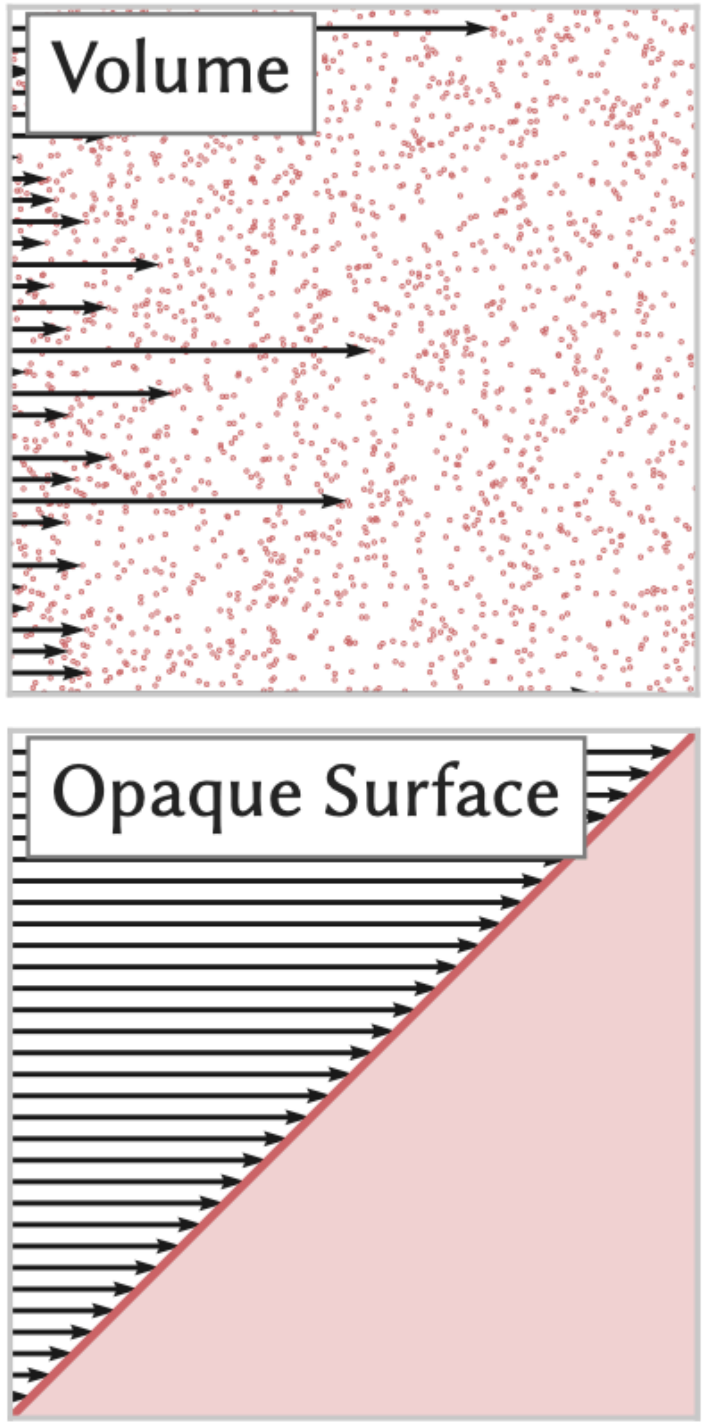

2D Scene

Volumetric scene representations

Transmittance due to different visibility

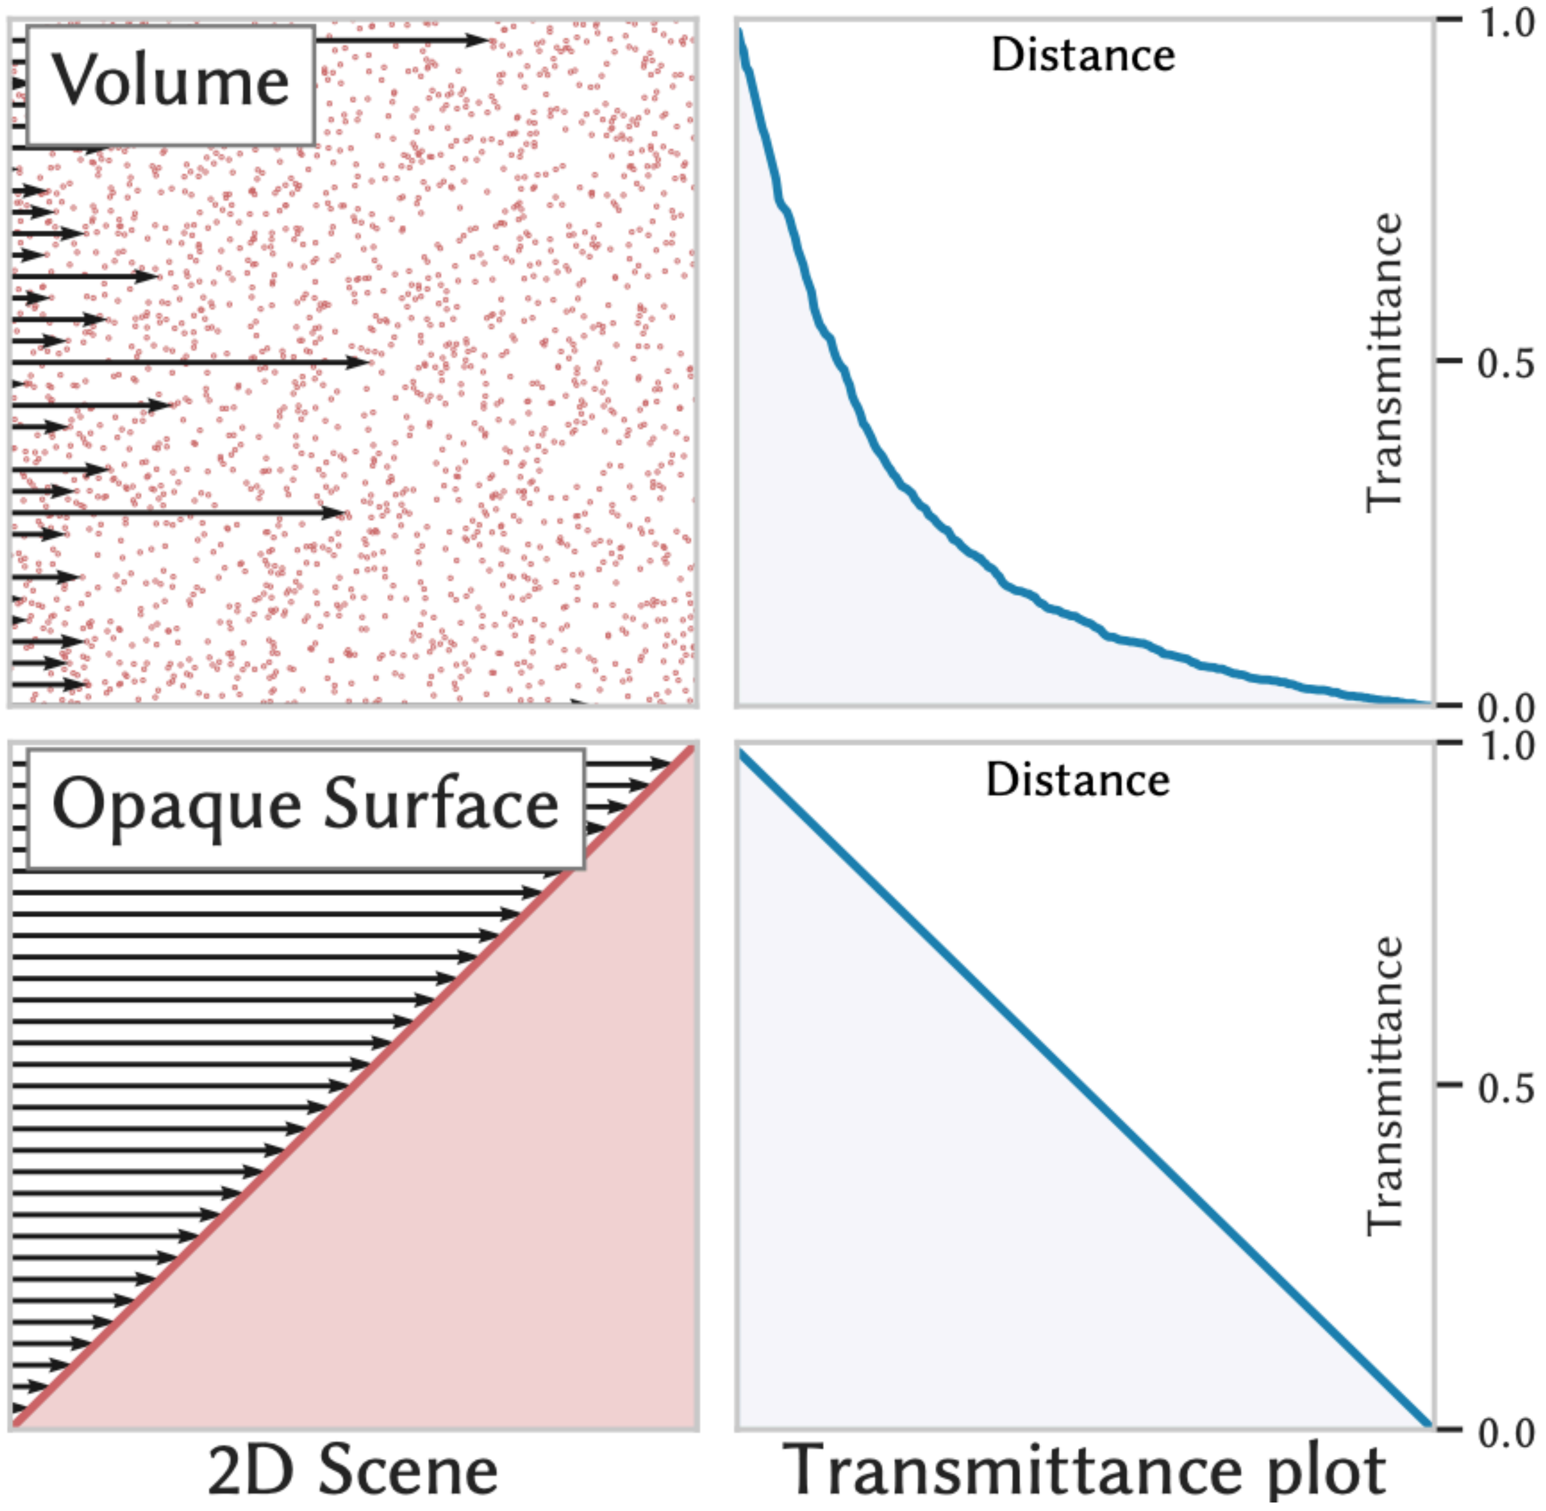

Volumetric scene representations

Transmittance due to different visibility

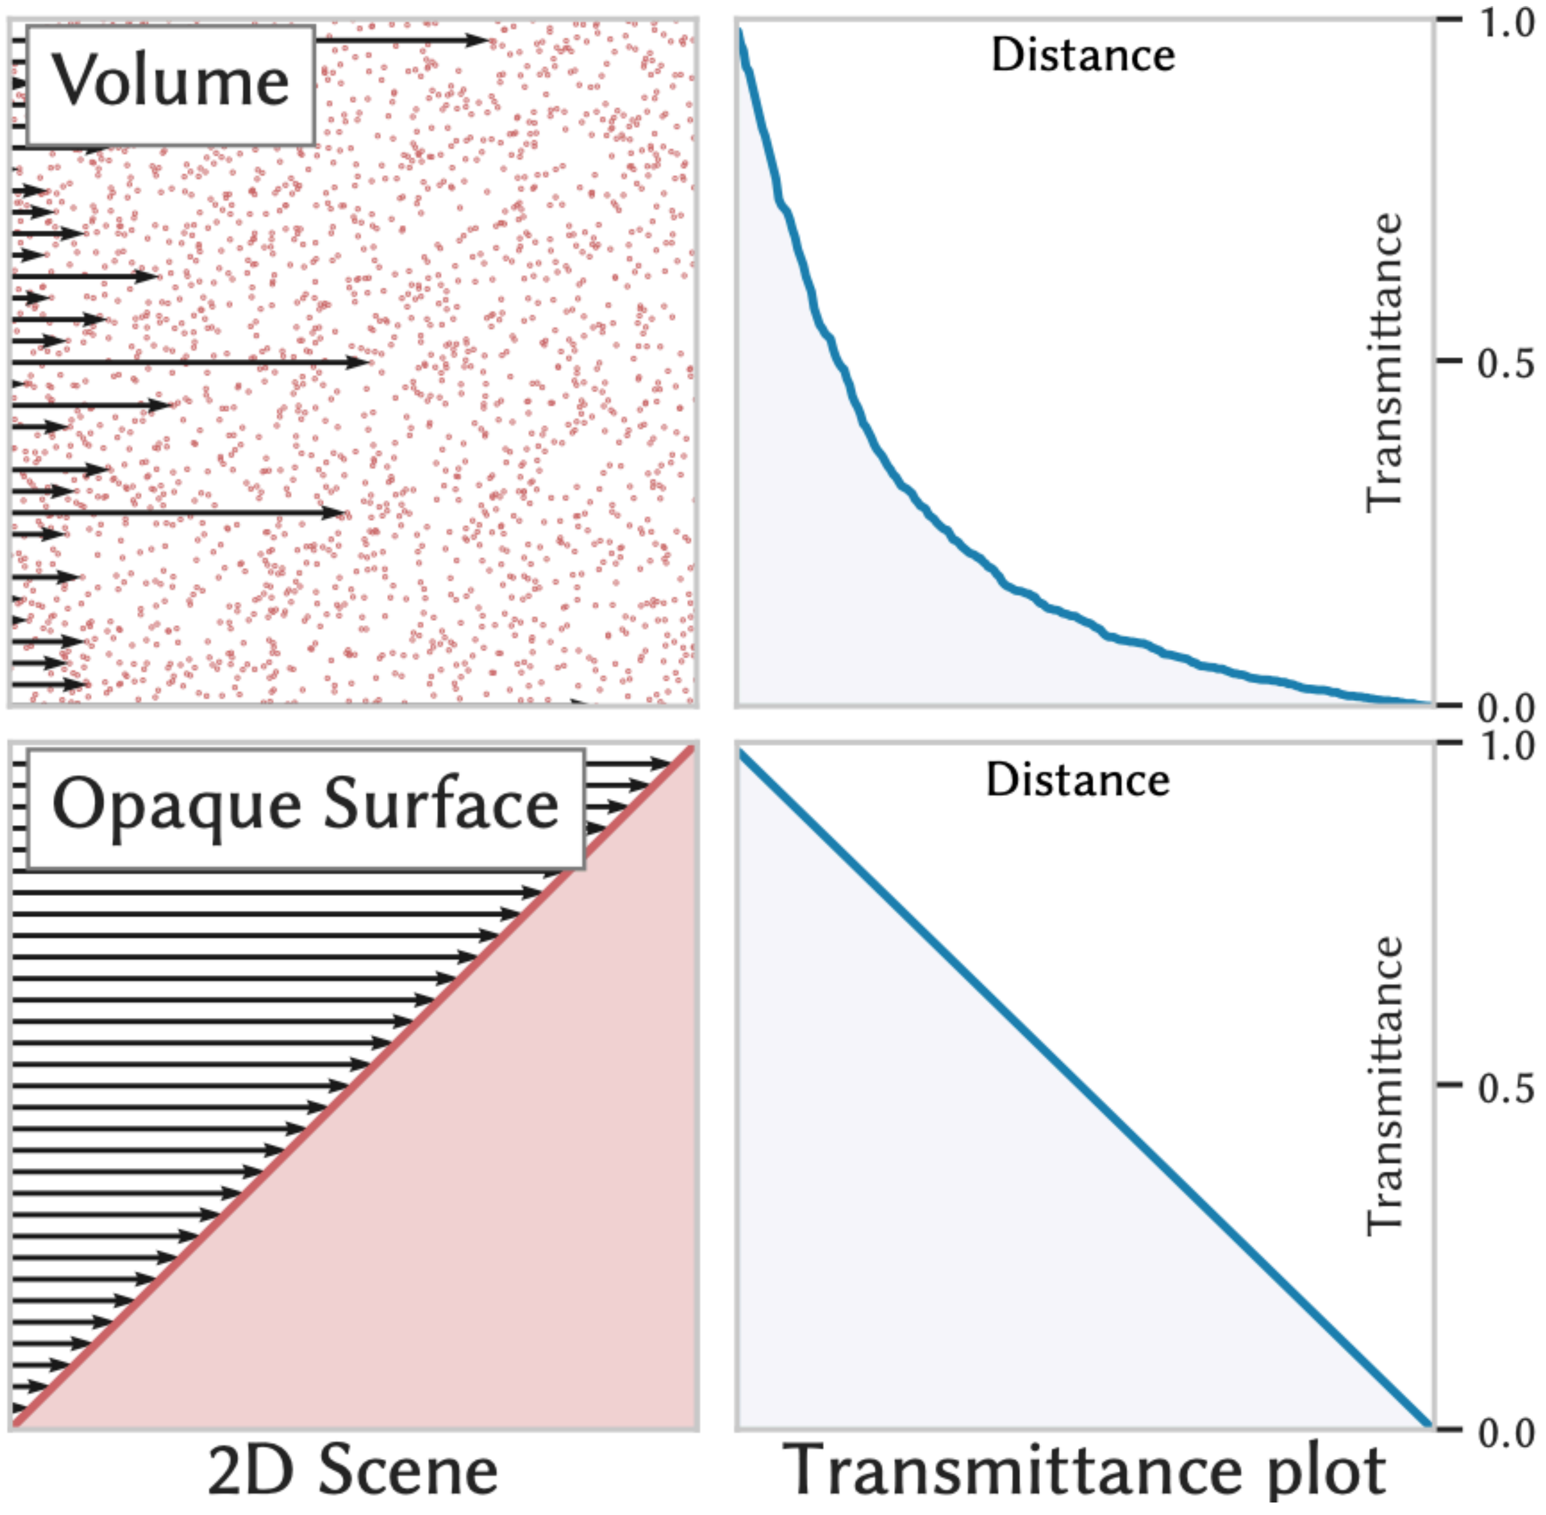

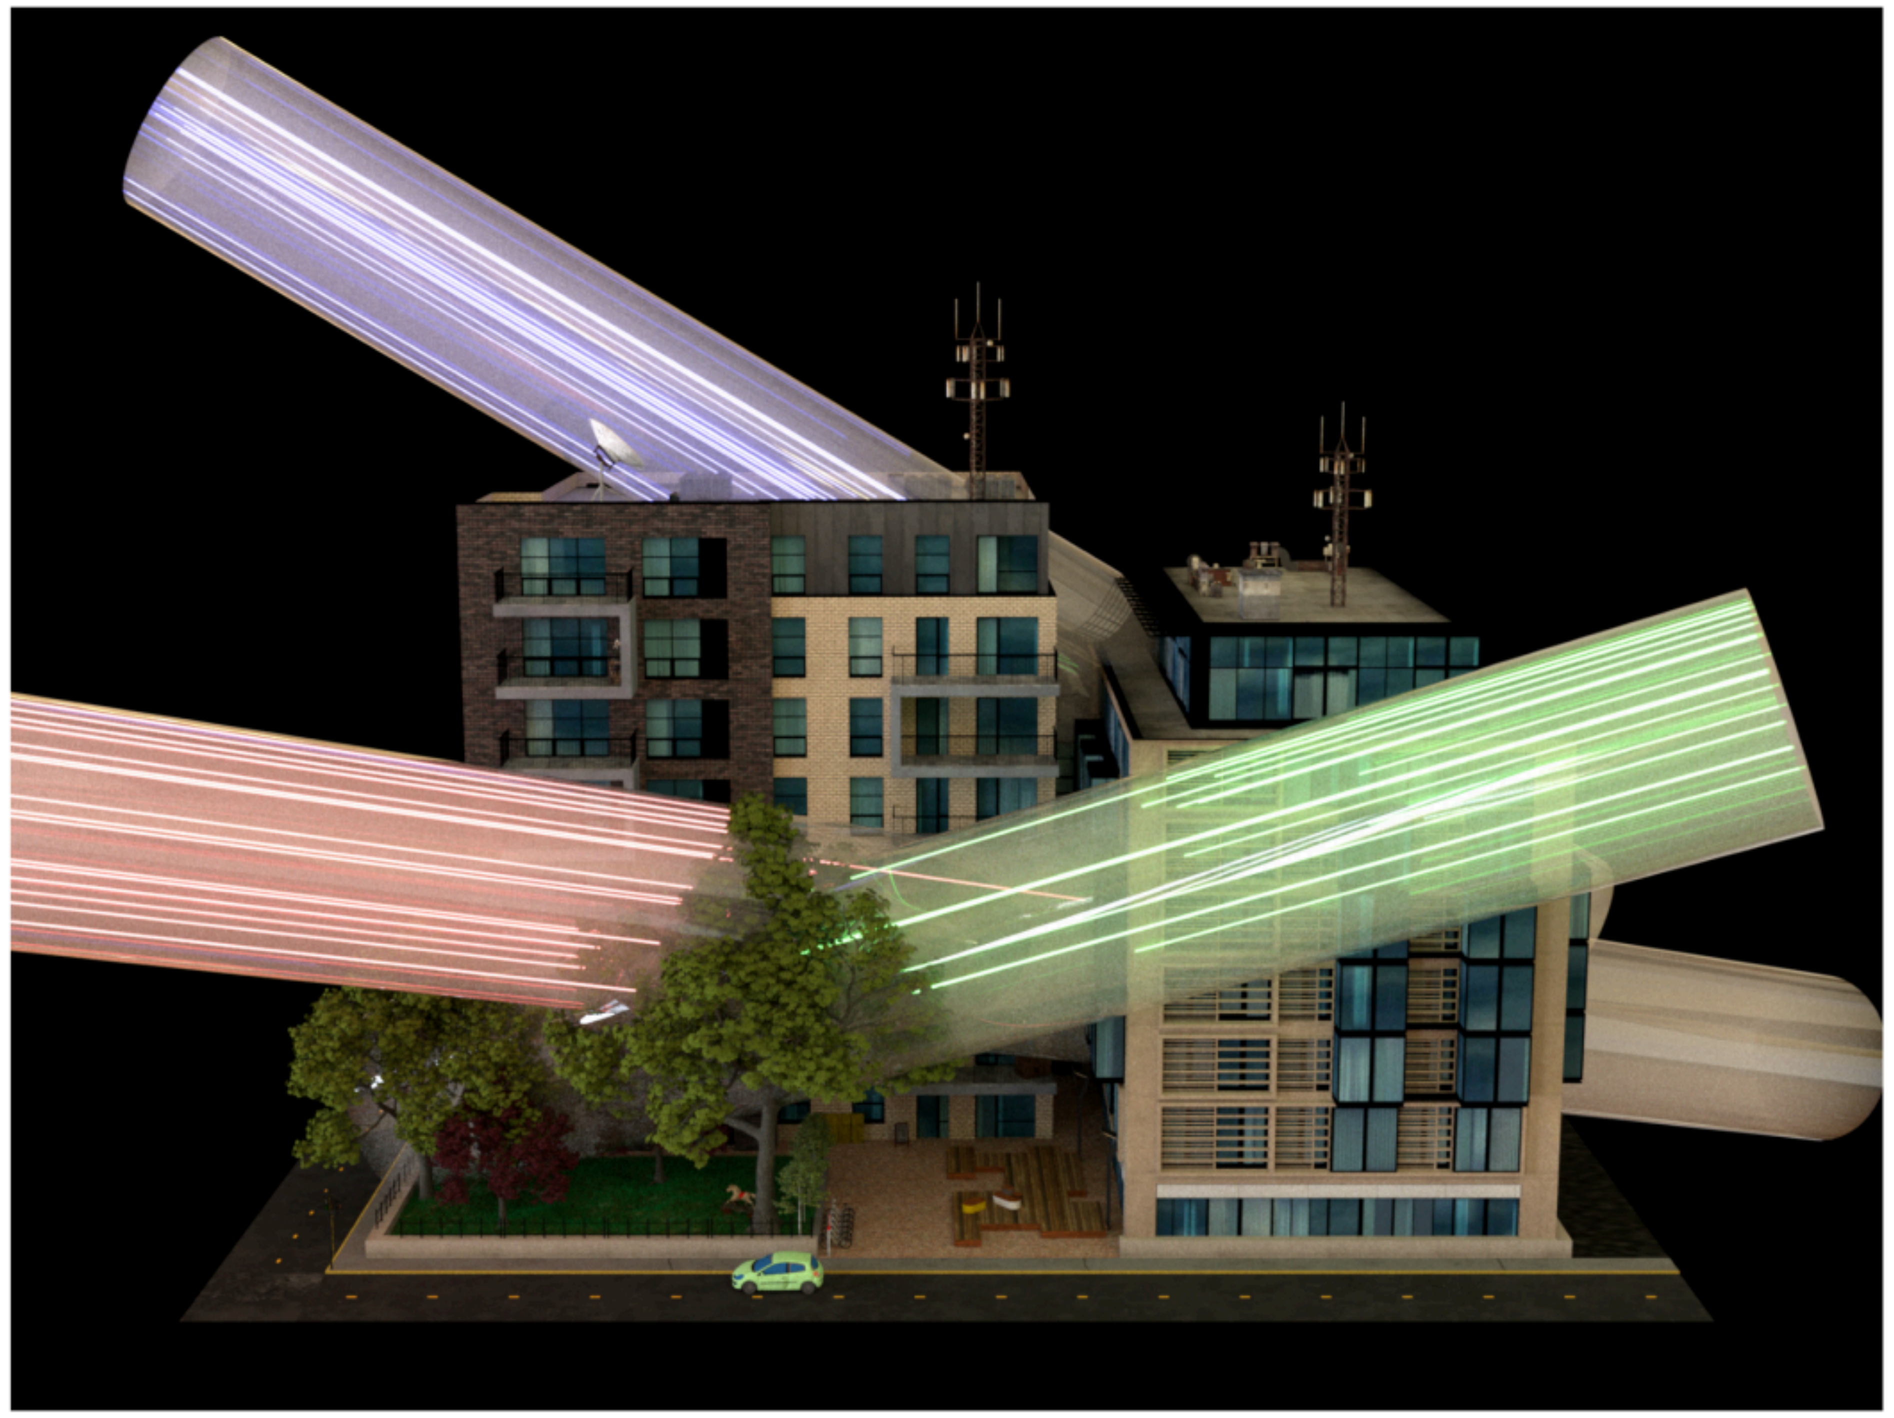

Volumetric scene representations

Transmittance due to different visibility

Volumetric scene representations

Basics

Volumetric scene representation has gain a lot of momentum for inverse rendering (NeRF 2020)

Volumetric representation results in more convex optimisation problem than directly optimising surface geometries

Volumetric representation is smooth and, unlike surface rendering, does not require any special treatment for visibility derivatives

Volumetric scene representations

Applications

Prefiltering for level of detail

Scene reconstruction using differentiable rendering

Neural rendering (NeRF)

Volumetric scene representations

Applications

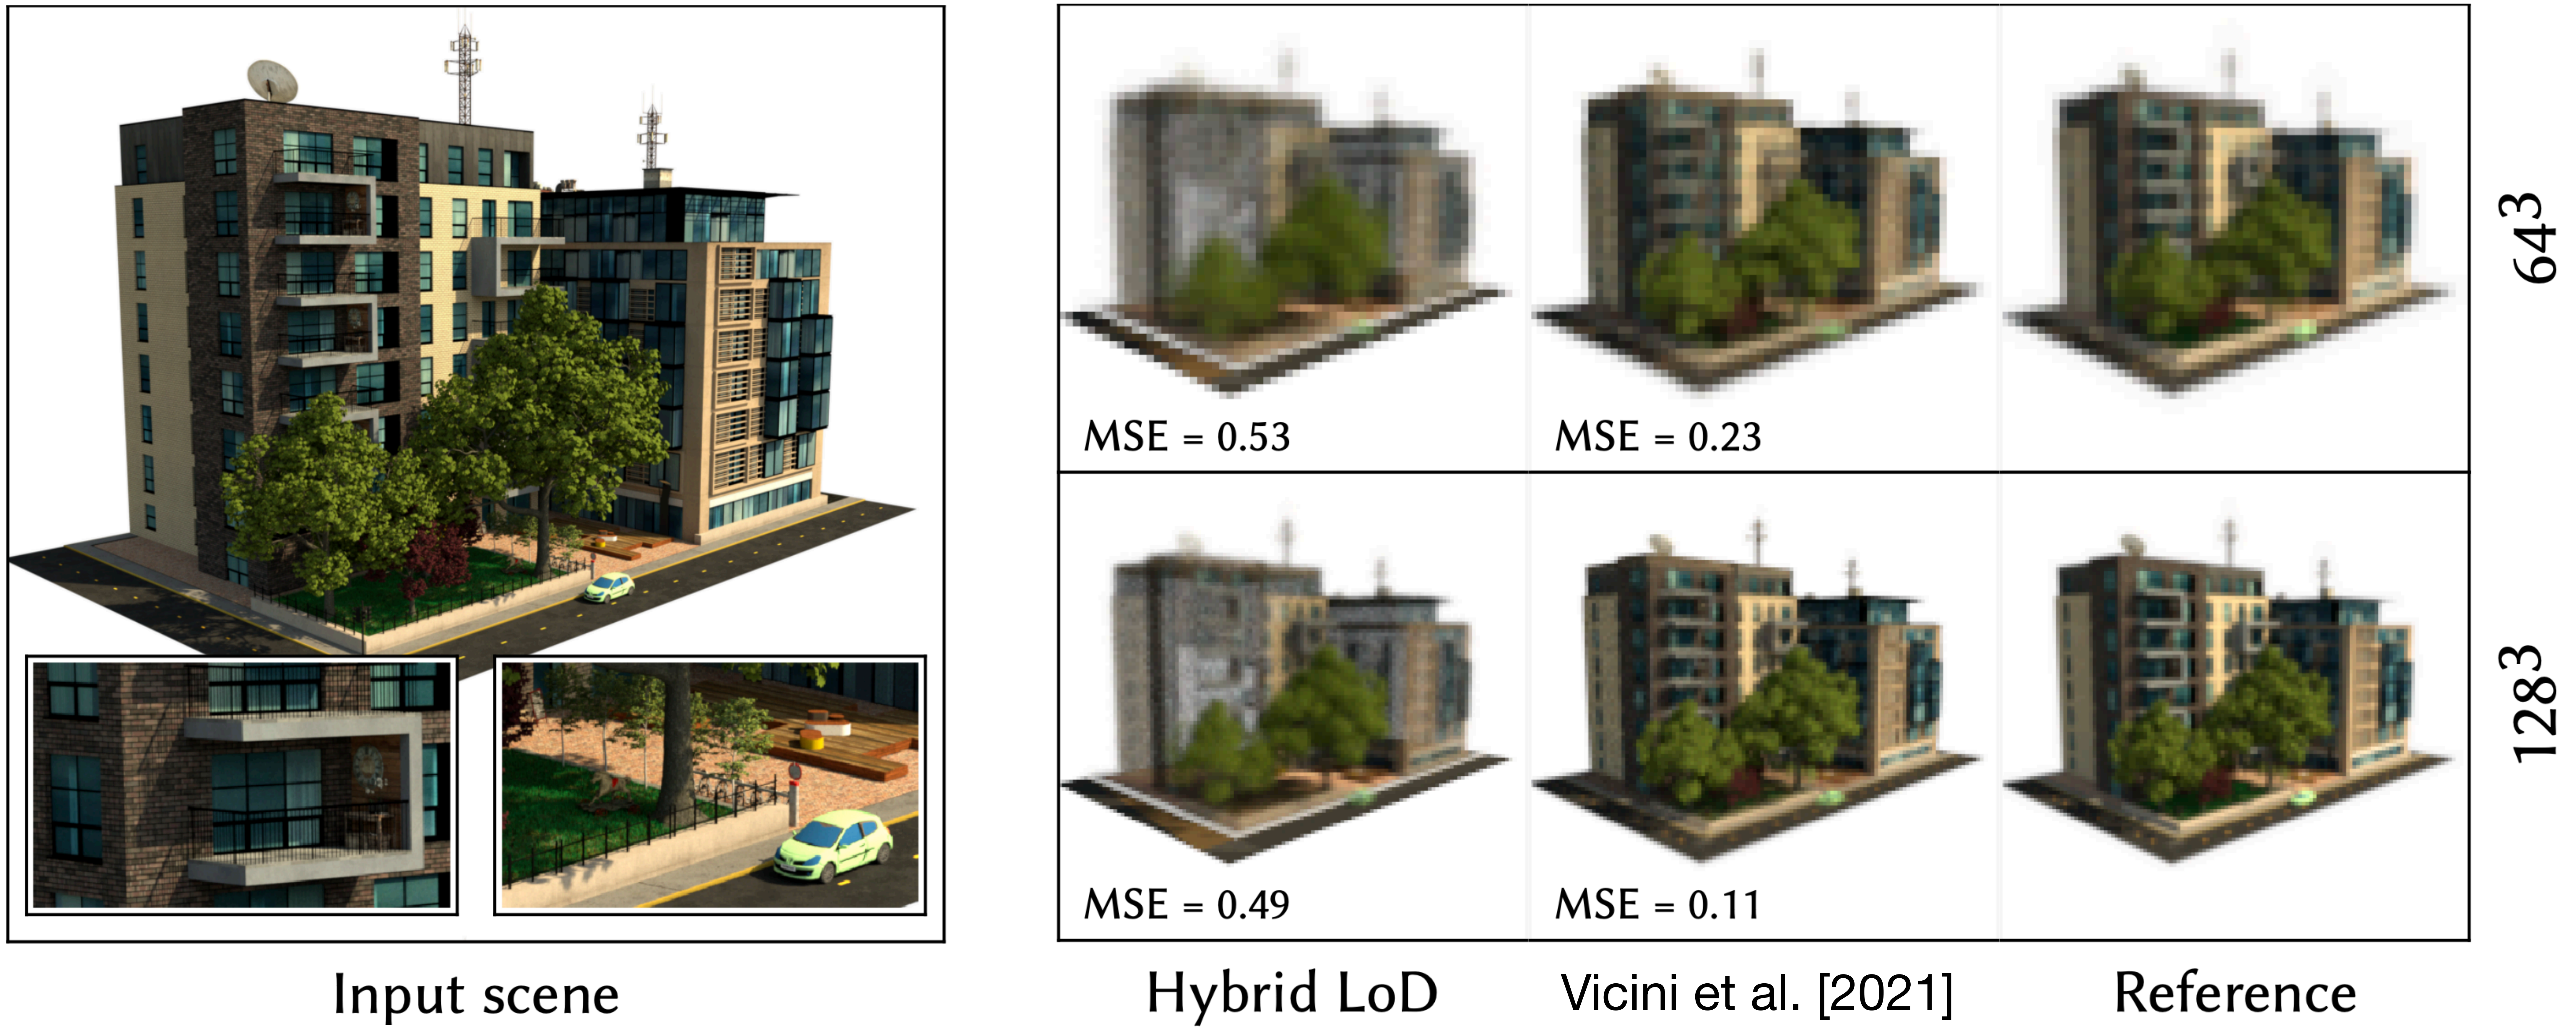

Neural scene representations for shading inference

Neural scene representation for shading inference

Neural scene representation for shading inference

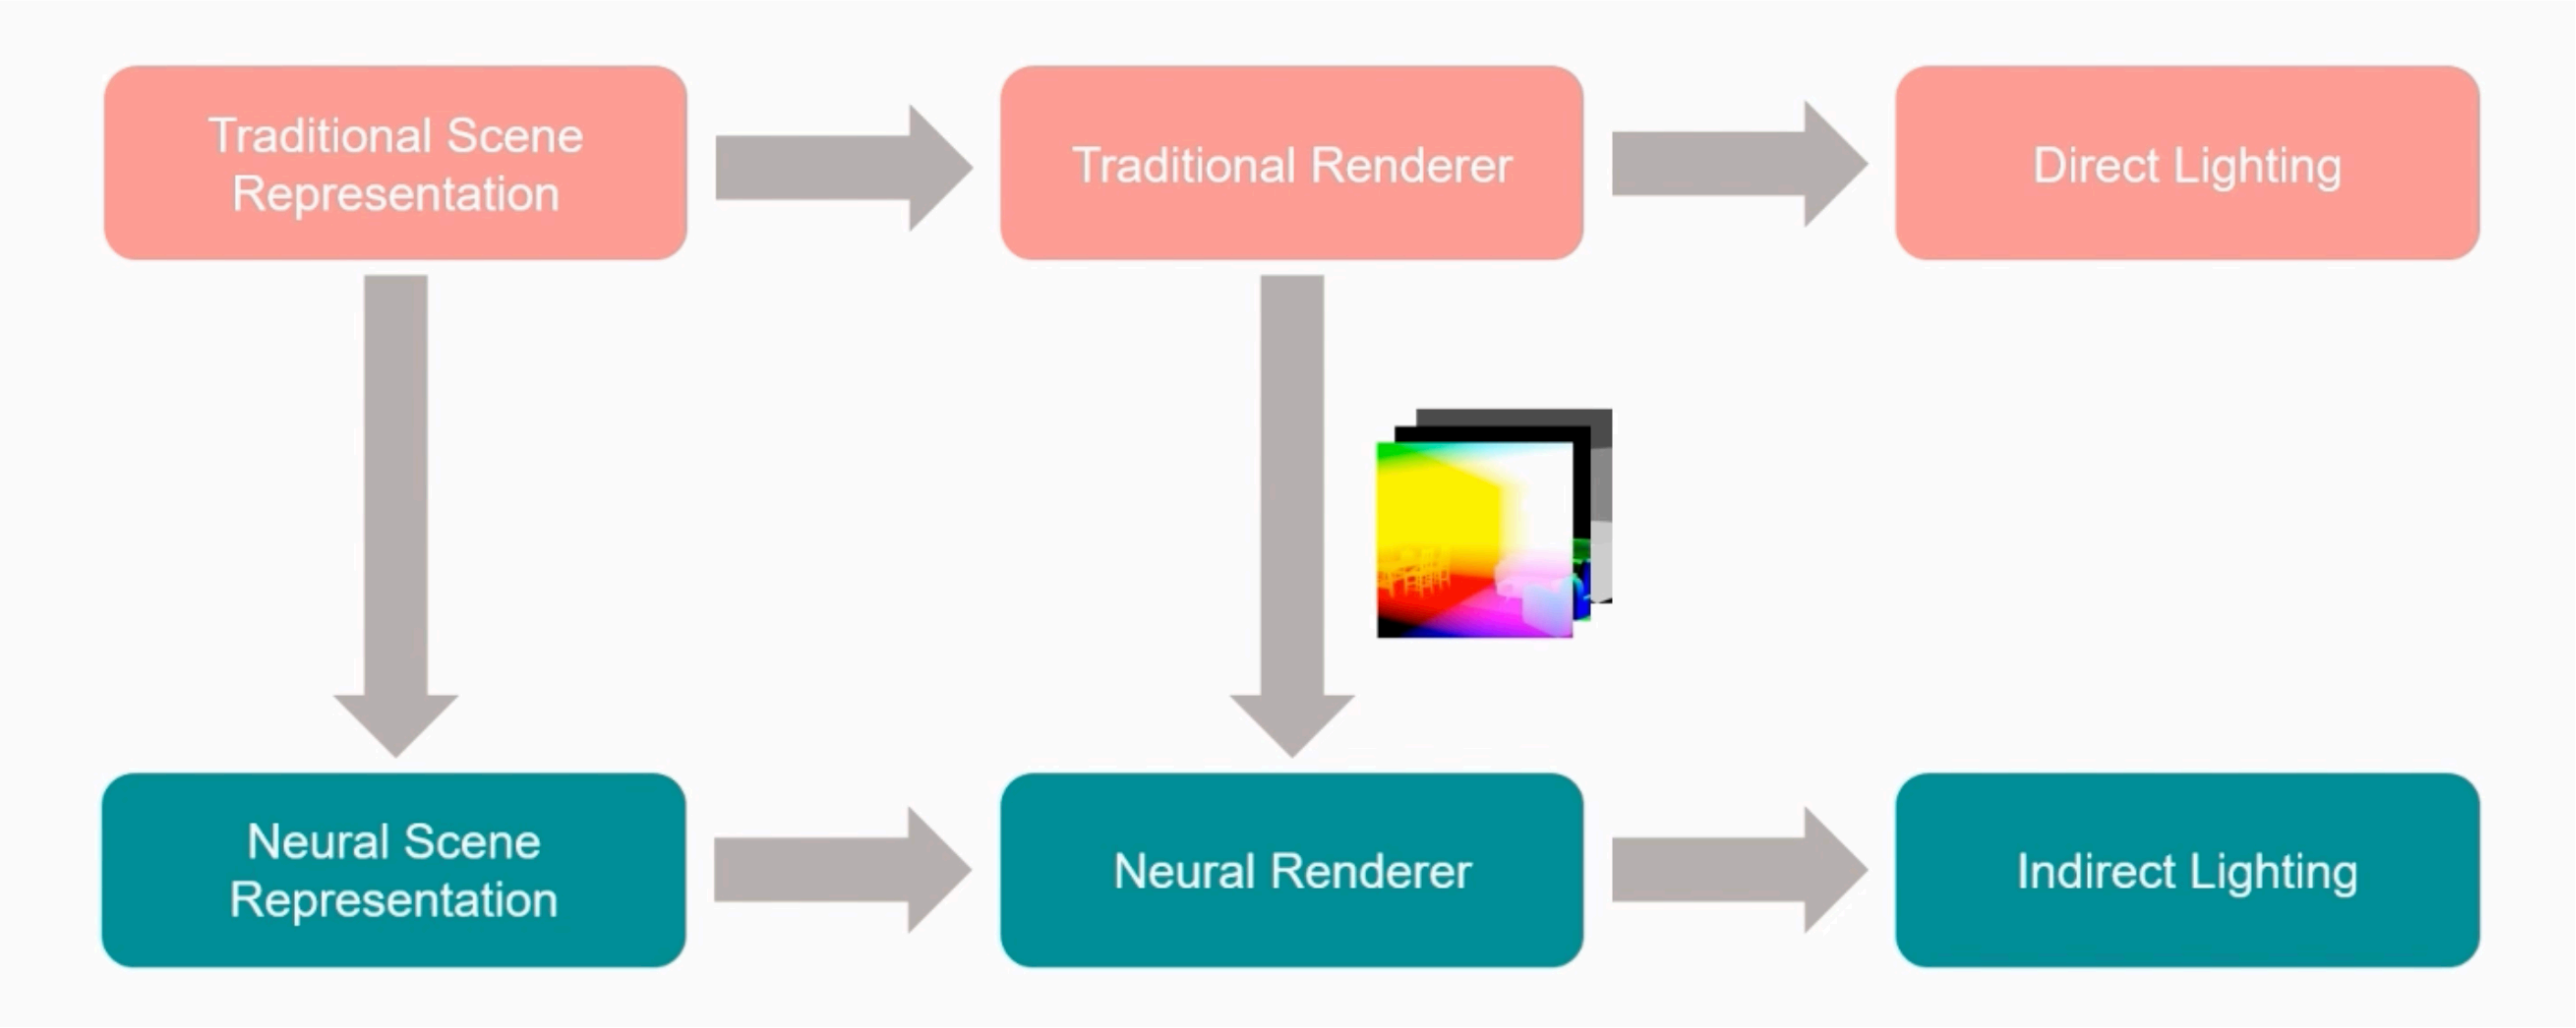

neural rendering

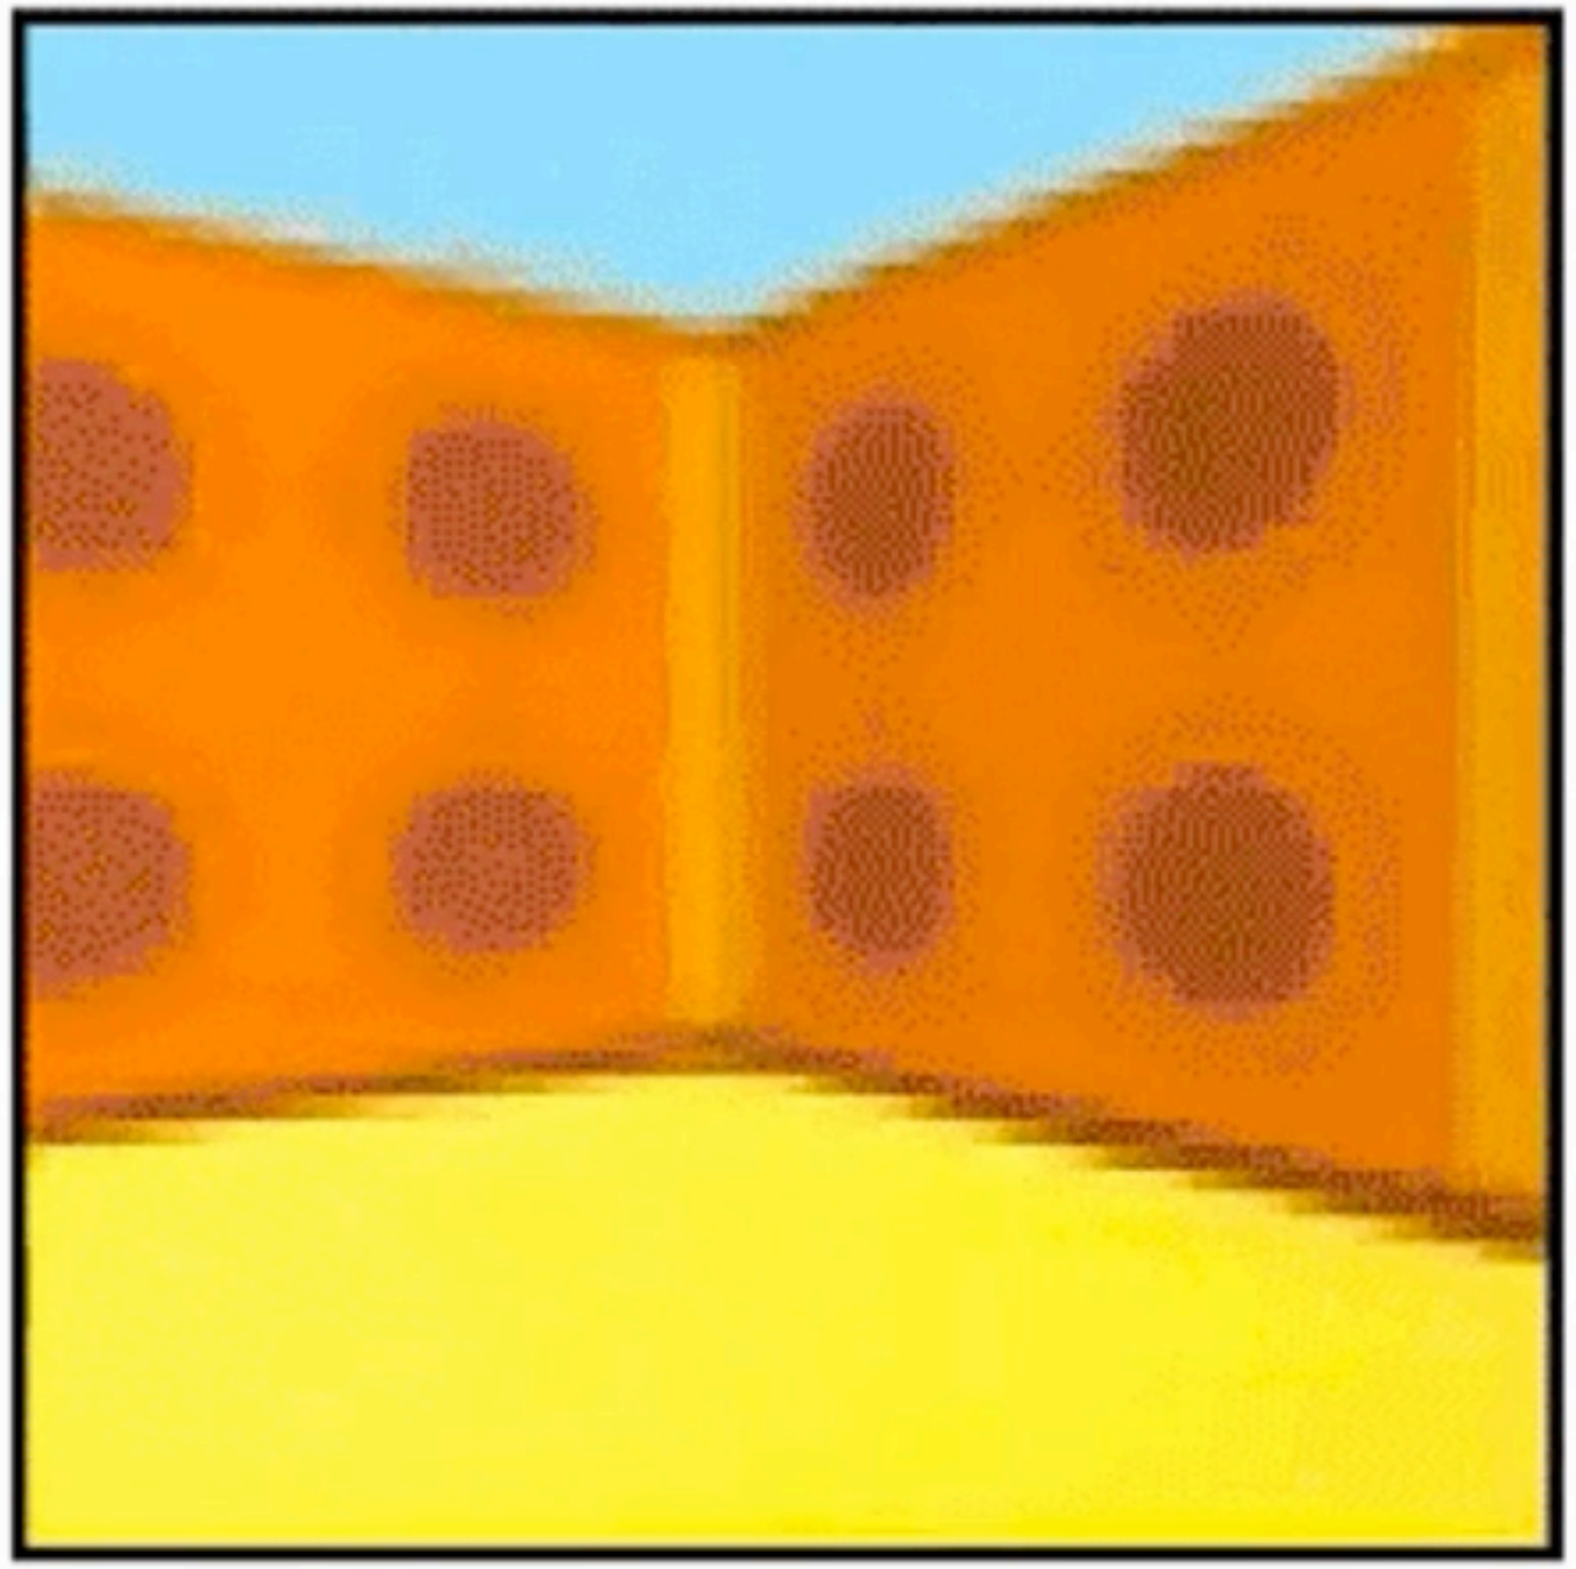

observation

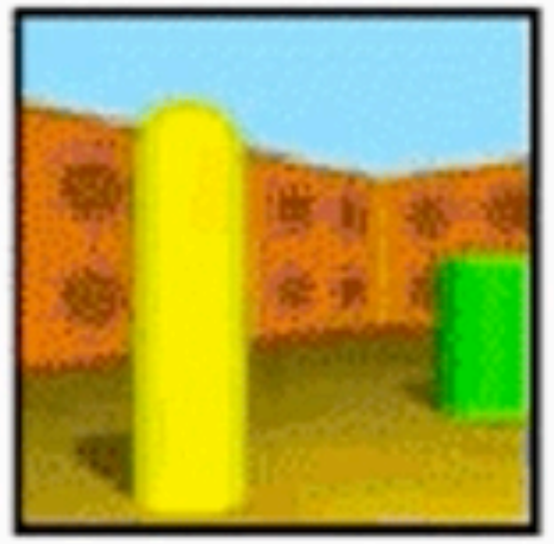

observation

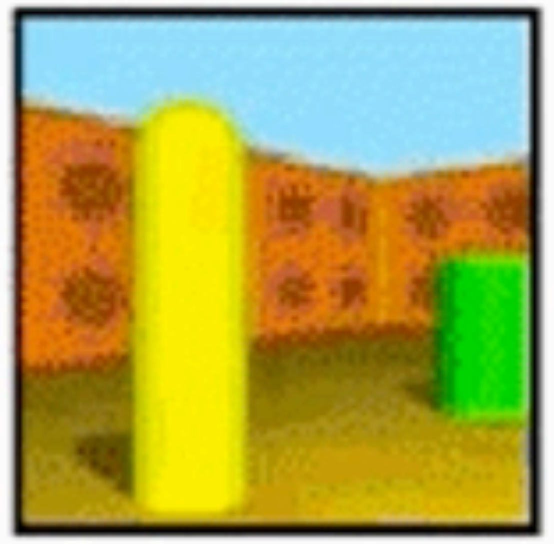

neural rendering

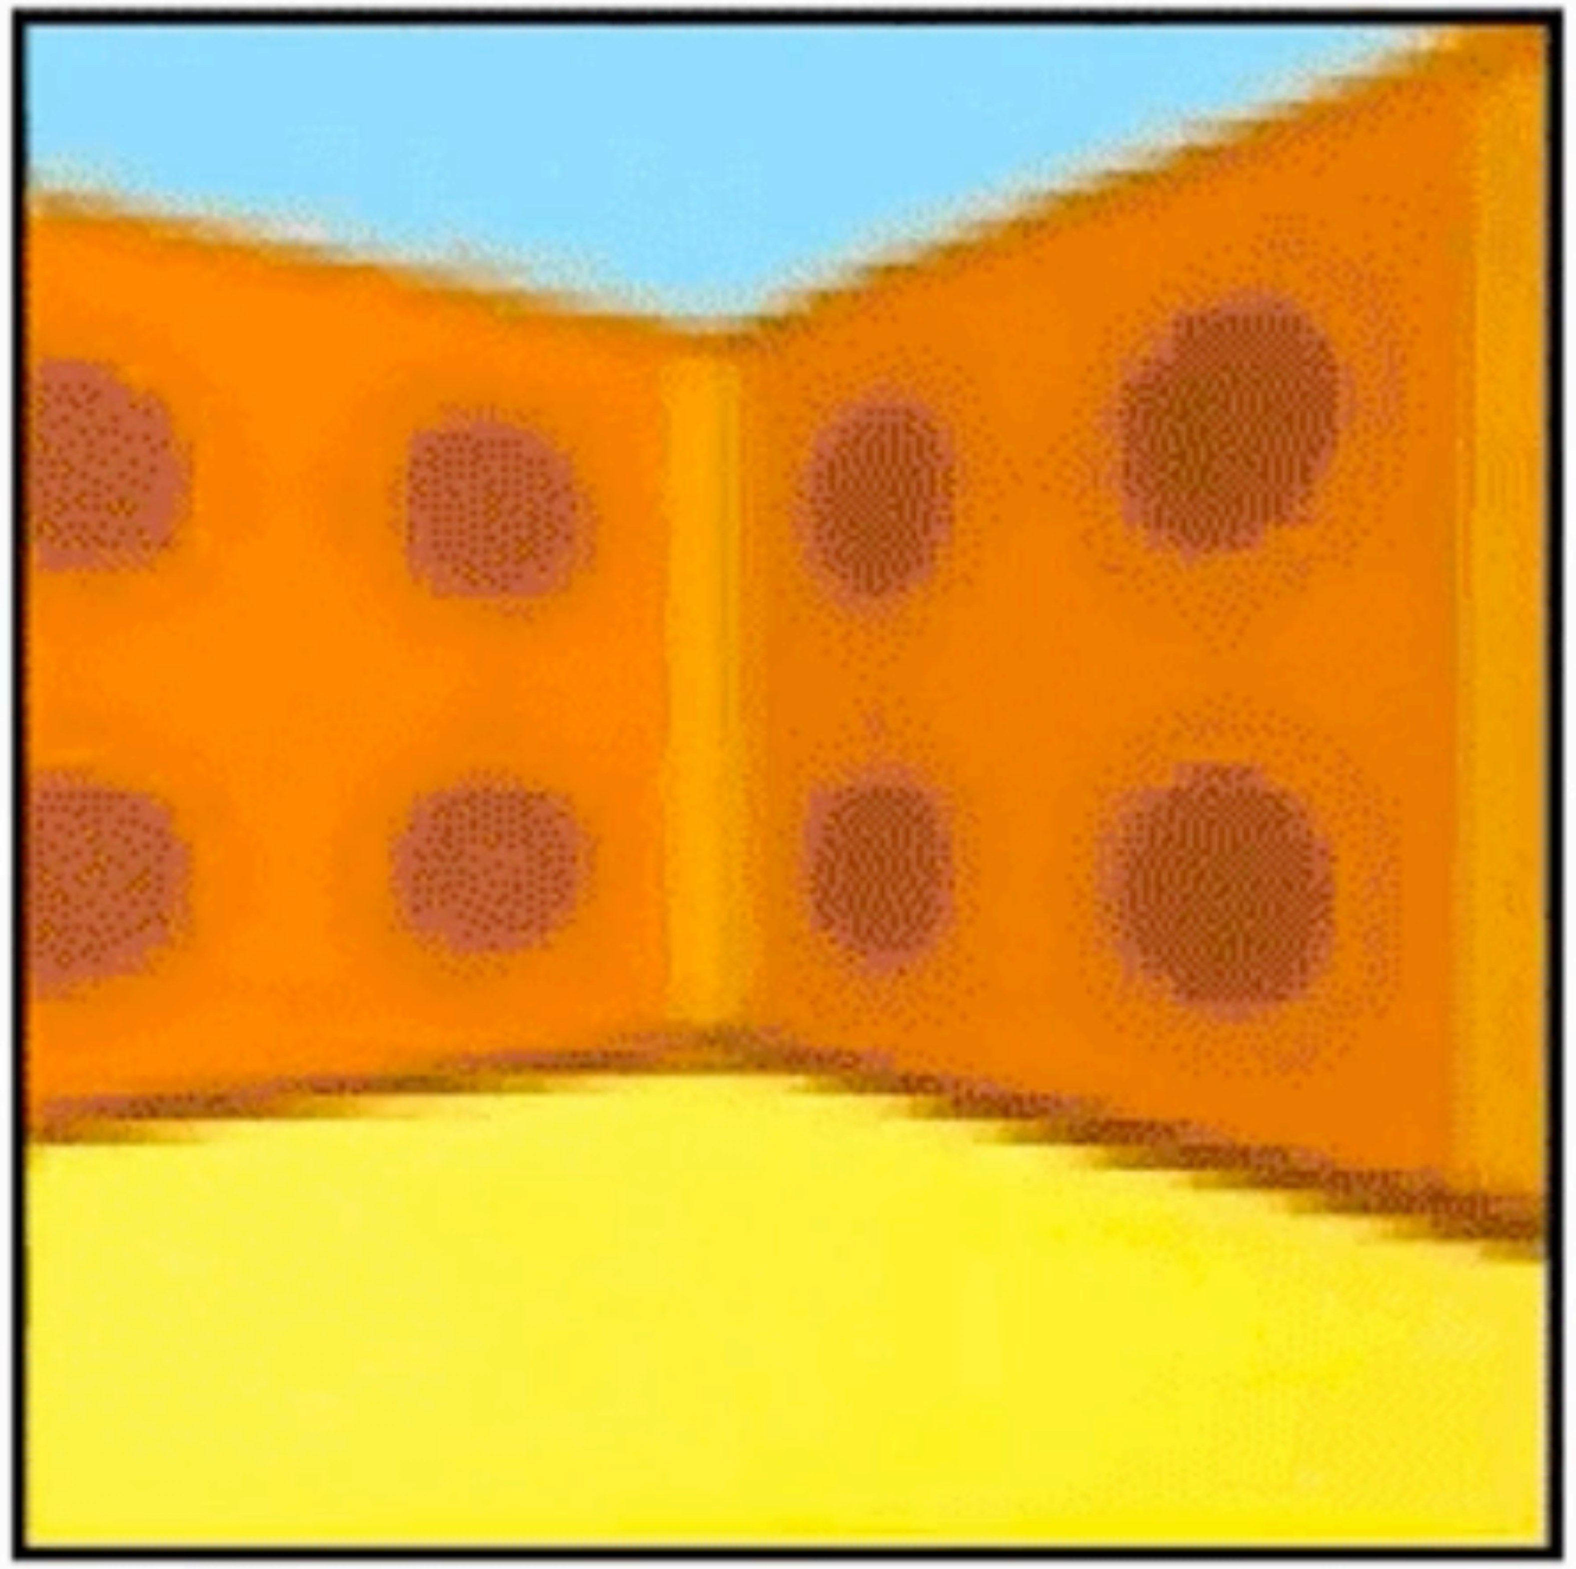

Neural scene representation for shading inference

References

Global Illumination with Radiance Regression Functions SIGGRAPH 2015

NeRF ECCV 2020

Compositional neural scene representations for shading inference SIGGRAPH 2020

Neural Control Variates SIGGRAPH Asia 2020

A Non-Exponential Transmittance Model for Volumetric Scene Representations SIGGRAPH 2021

Weakly Supervised Contrastive Learning in Path Manifold for Monte Carlo Image Reconstruction SIGGRAPH 2021

Acknowledgements

Special thanks to

Muller et al. [2019,2020] to make their slides and videos available online

Granskog et al. [2020] and Mildenhall et al. [2020] for making their video presentations available online

Vicini et al. [2021] and Cho et al. [2021] for releasing the preprint of their recent work before publication