# NeRF++: Analyzing and Improving Neural Radiance Fields

Kai Zhang
Cornell University
GAMES 2020/6/21

# What is view synthesis?

28 captured photos

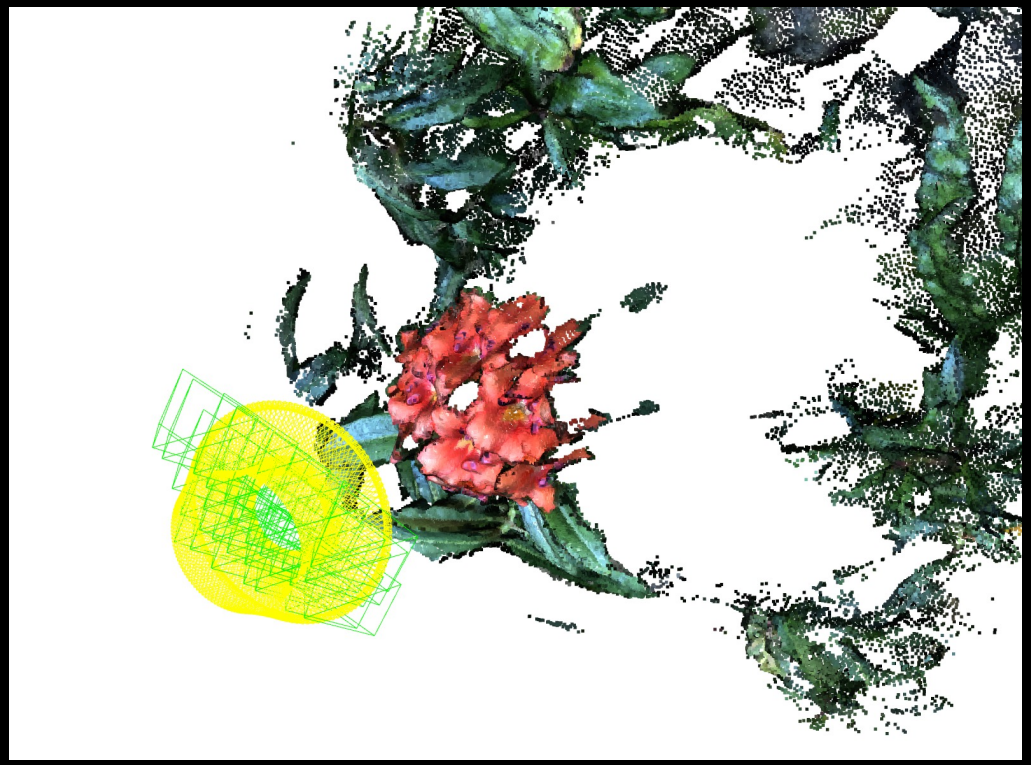

Given captured green camera views, synthesize smooth yellow camera trajectory views

Mildenhall et al., 2020

# NeRF: Neural Radiance Fields

Mildenhall et al., 2020

Mildenhall et al., 2020

# How NeRF works?

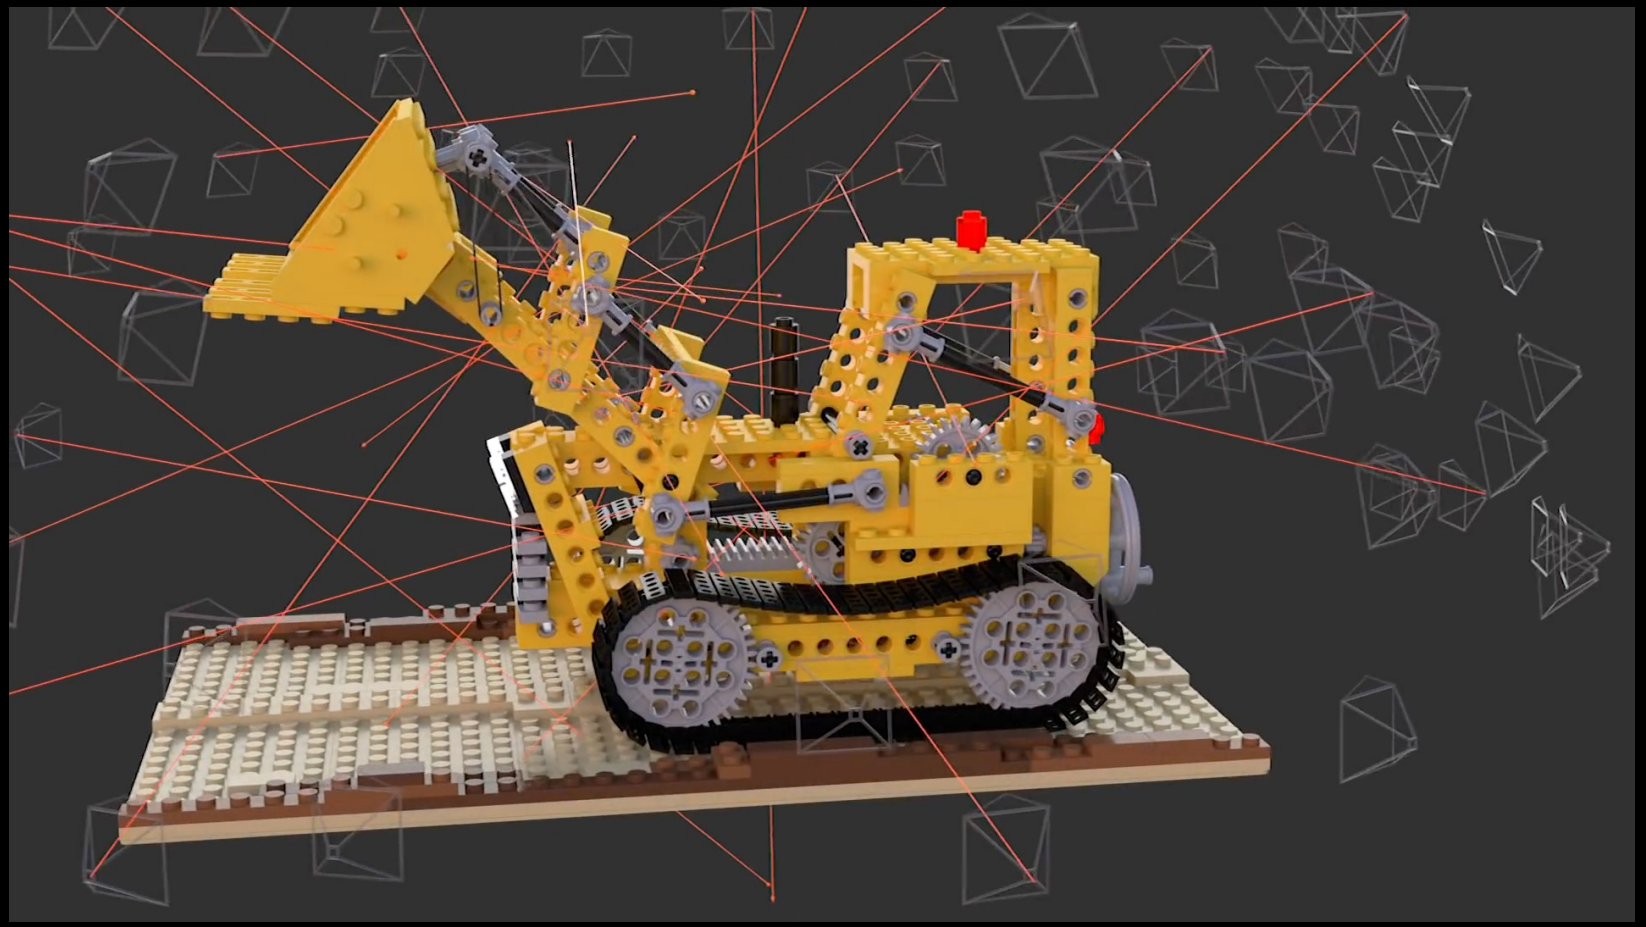

Mildenhall et al., 2020

# Break NeRF into geometry and appearance modelling

## Geometry

Soft opacity field (fog)

$$(x, y, z) \rightarrow \sigma$$

## Appearance

Radiance field (surface light field [1])

$$(x, y, z, \theta, \phi) \rightarrow rgb$$

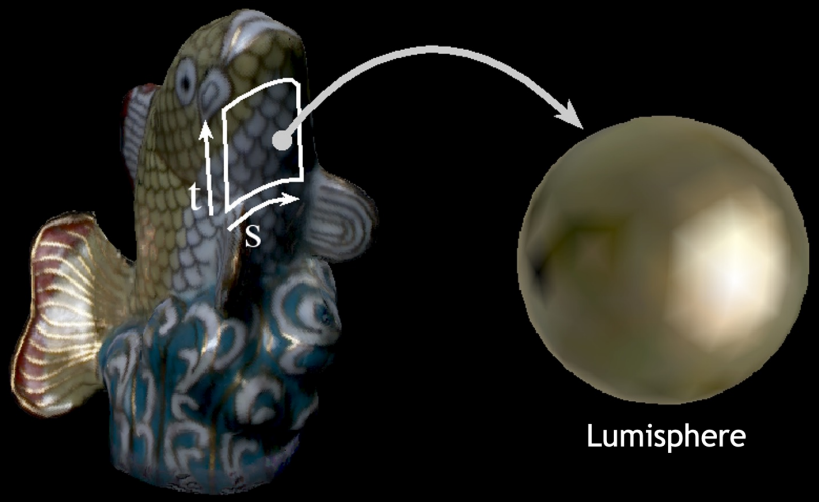


[1] Wood, Daniel N., et al. "Surface light fields for 3D photography." Siggraph 2000.

# Remarks on NeRF’s geometry modelling

- Soft geometry
  - initialized as nothingness
  - grow as needed

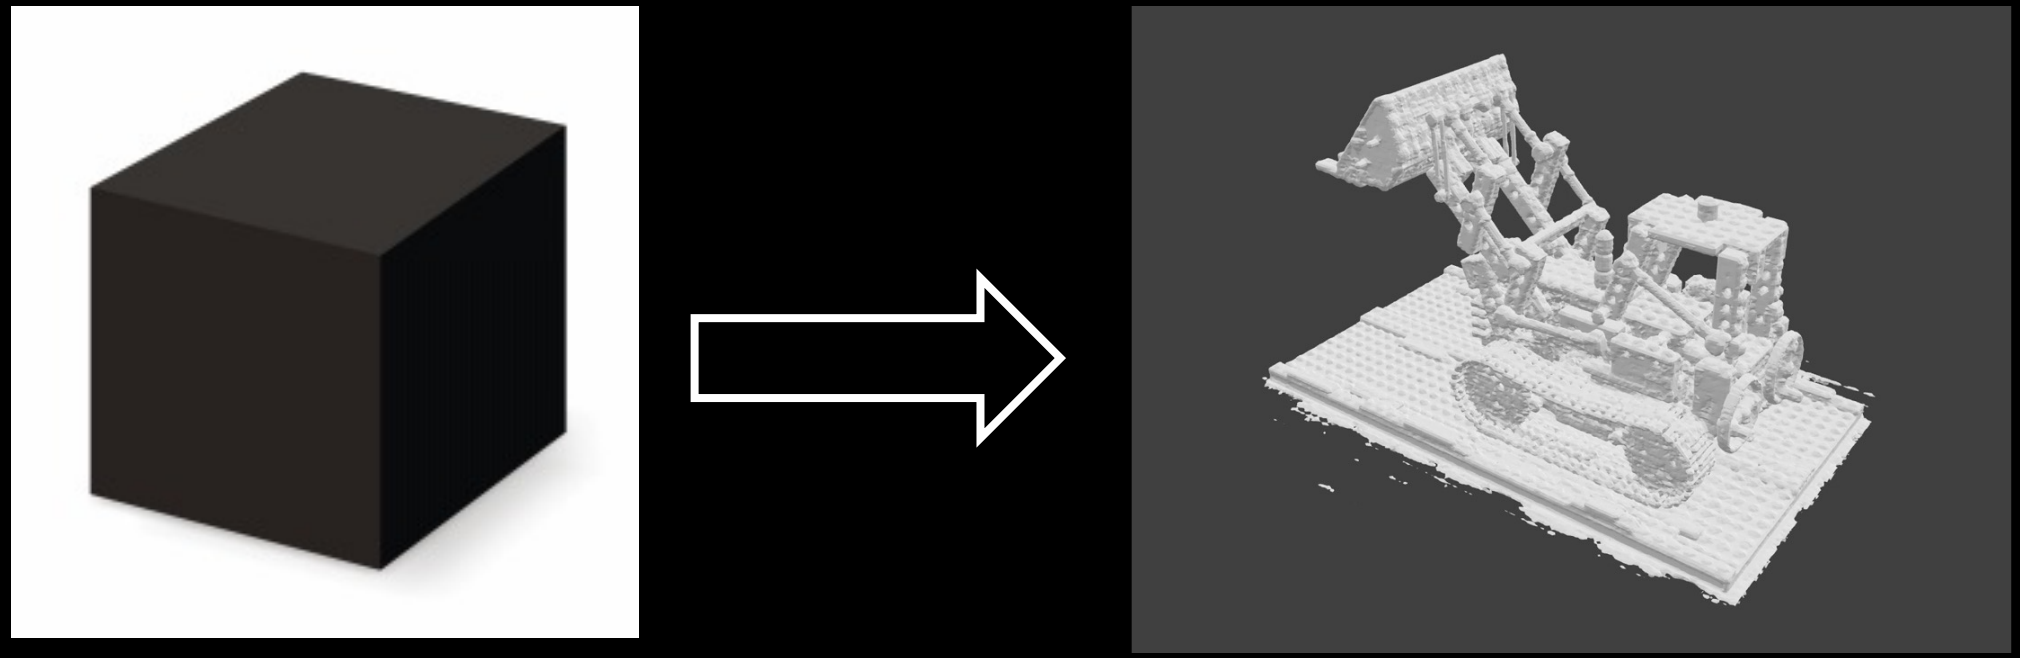

# Remarks on NeRF's geometry modelling

- Compare with works using hard geometry, e.g., DVR [1], IDR [2]
  - ✓ Not require object segmentation masks: genus issue
  - ✓ No boundary discontinuity: easy for differentiable rendering
  - ✓ Easy to optimize shape: more robust to shape-radiance ambiguity

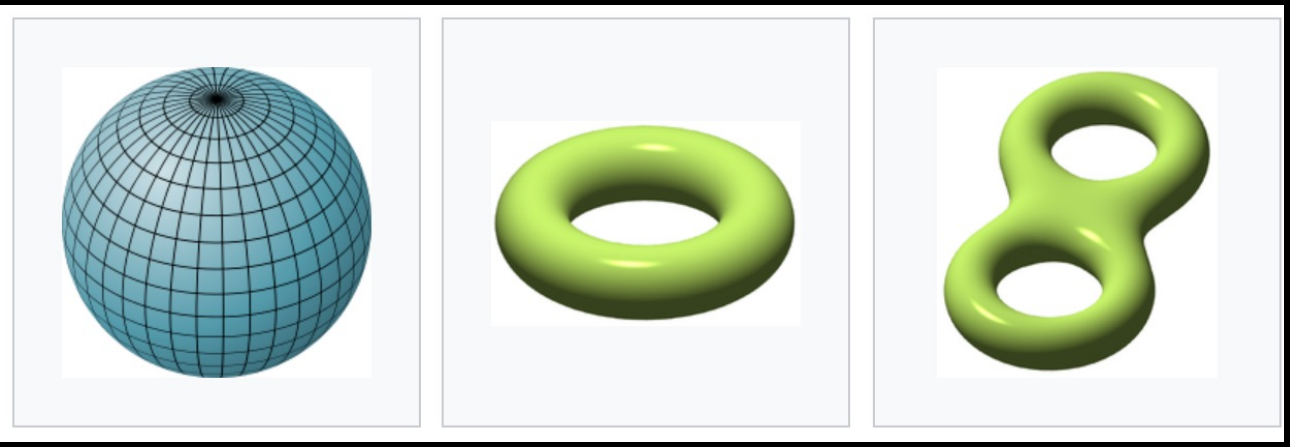

Genus 0 Genus 1 Genus 2

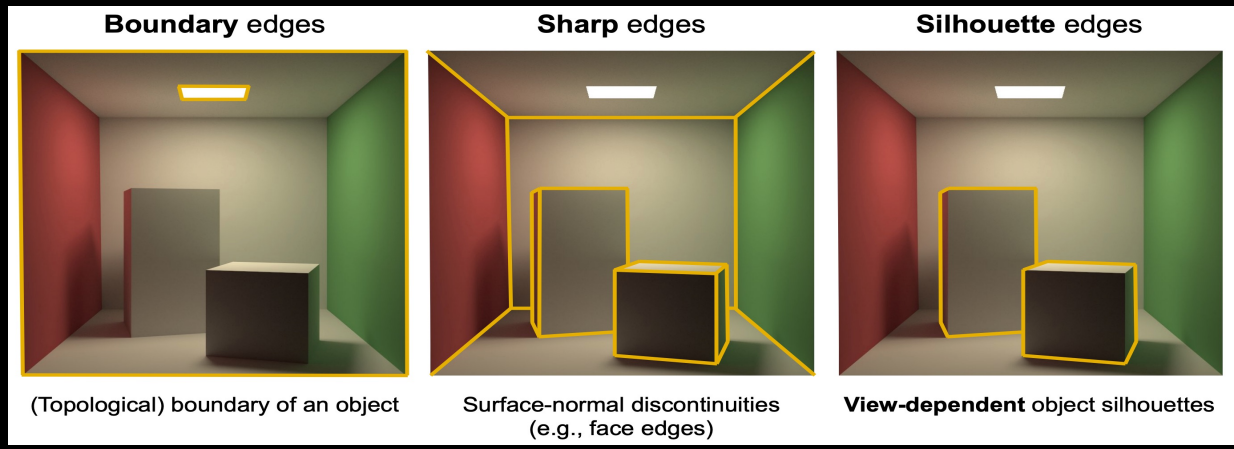


Zhao et al. Physics-based differentiable rendering


[1] Niemeyer, Michael, et al. "Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision." *CVPR* 2020.
[2] Yariv, Lior, et al. "Multiview neural surface reconstruction by disentangling geometry and appearance." *NeuralIPS* 2020.

# Shape-radiance ambiguity

- Infinitely many shape-radiance combinations can explain training views.
- But not all of them generalize to novel viewpoints.

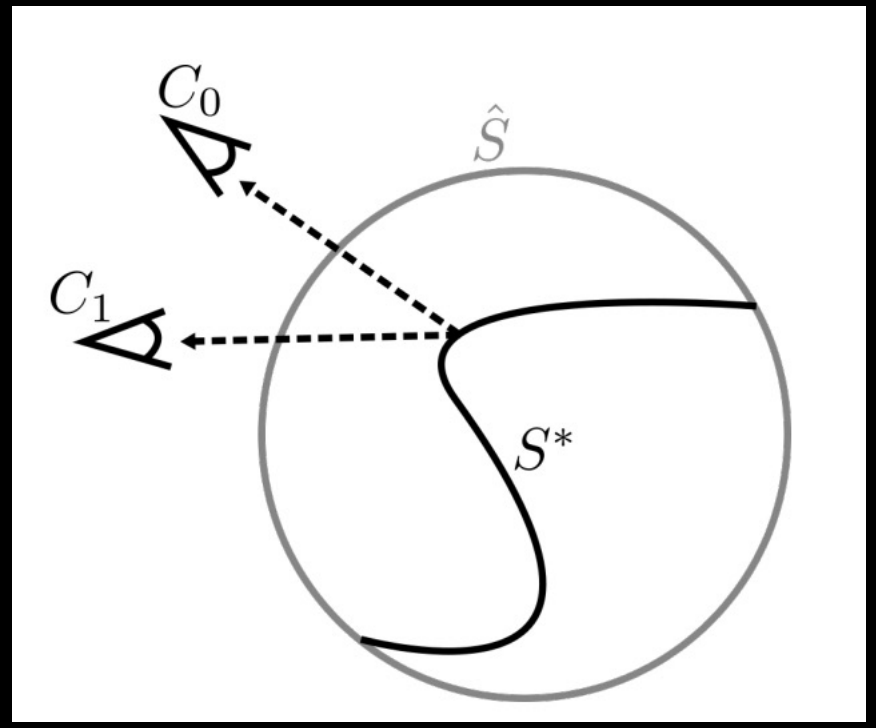


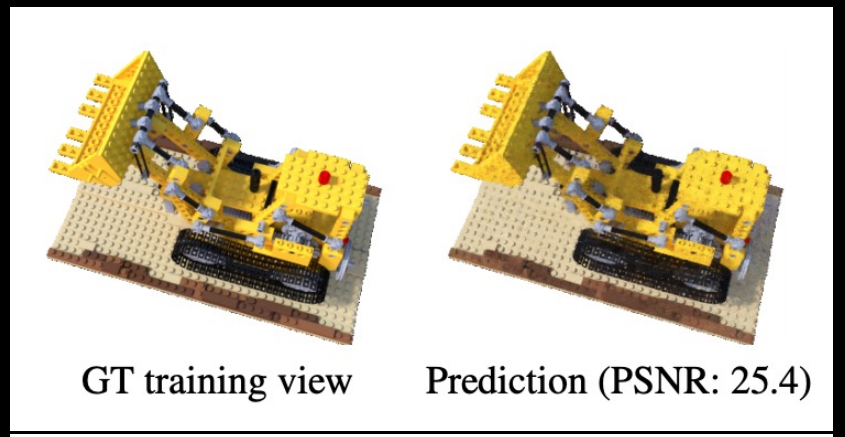


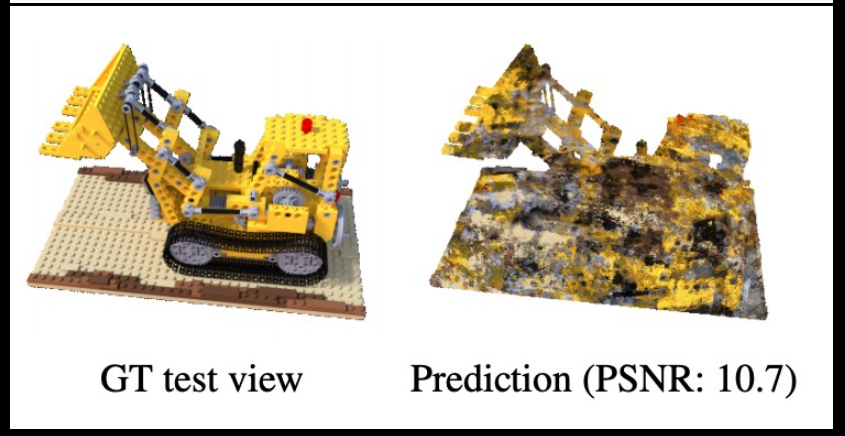

# Shape-radiance ambiguity

Why doesn’t NeRF fall into such trivial solutions?

- Factor 1: easily optimizable soft geometry

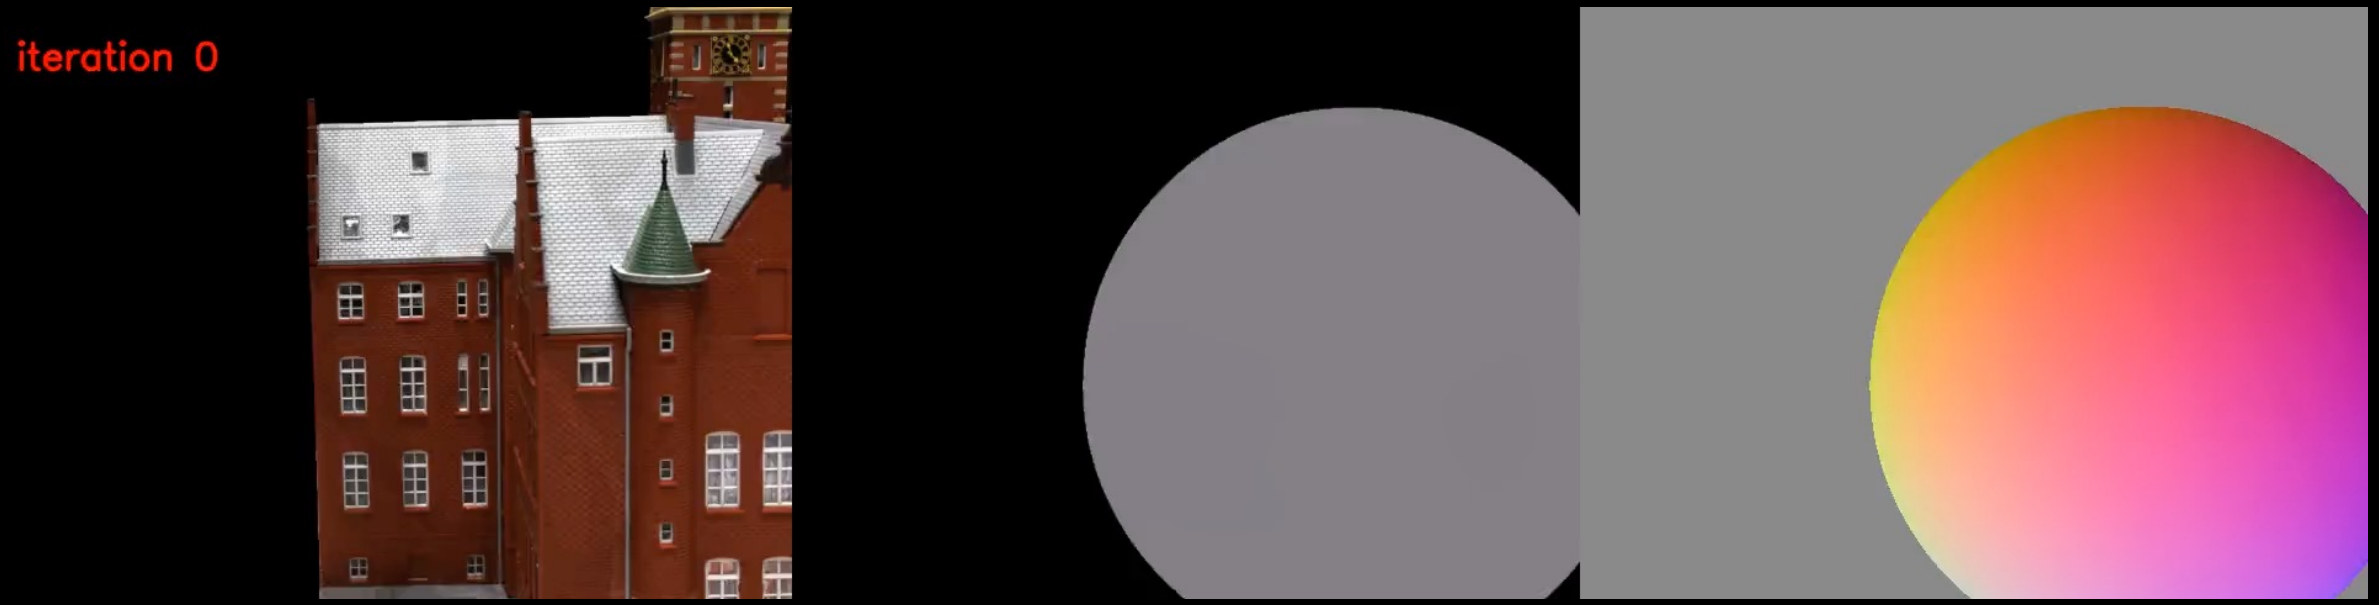


Switch to a hard geometry representation: SDF

# Shape-radiance ambiguity

Why doesn’t NeRF fall into such trivial solutions?

- Factor 2: smooth BRDF prior.

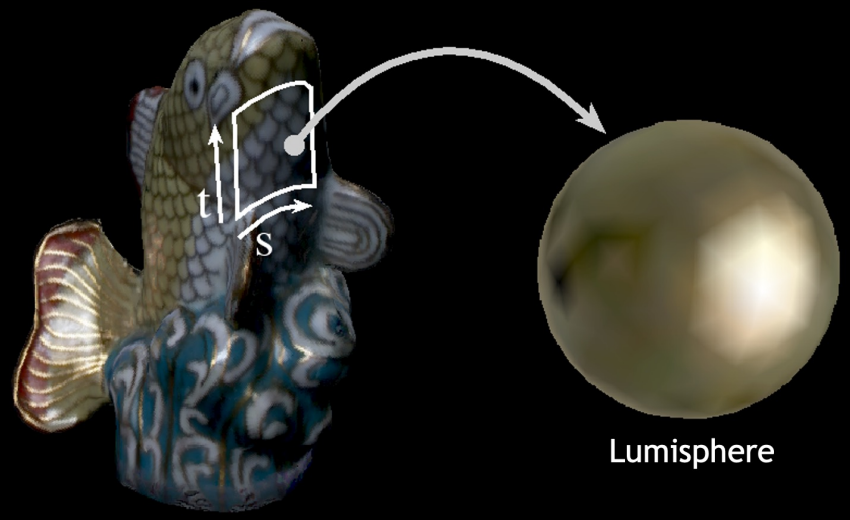


Prior in the data

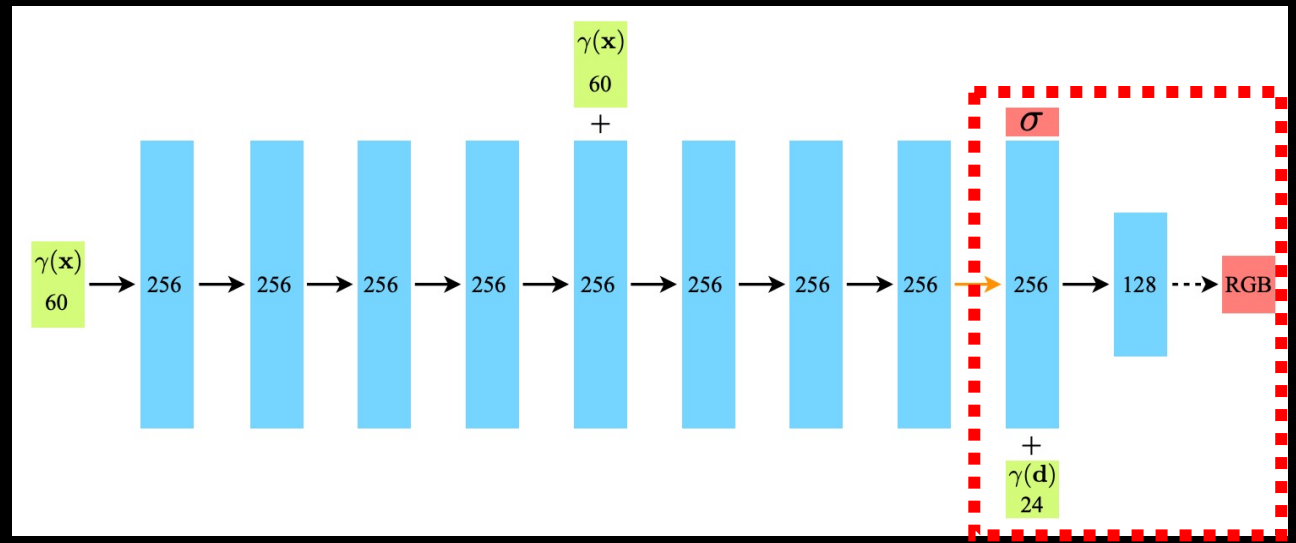


Prior in the algorithm

# Shape-radiance ambiguity

Why doesn't NeRF fall into such trivial solutions?

- Factor 2: smooth BRDF prior.

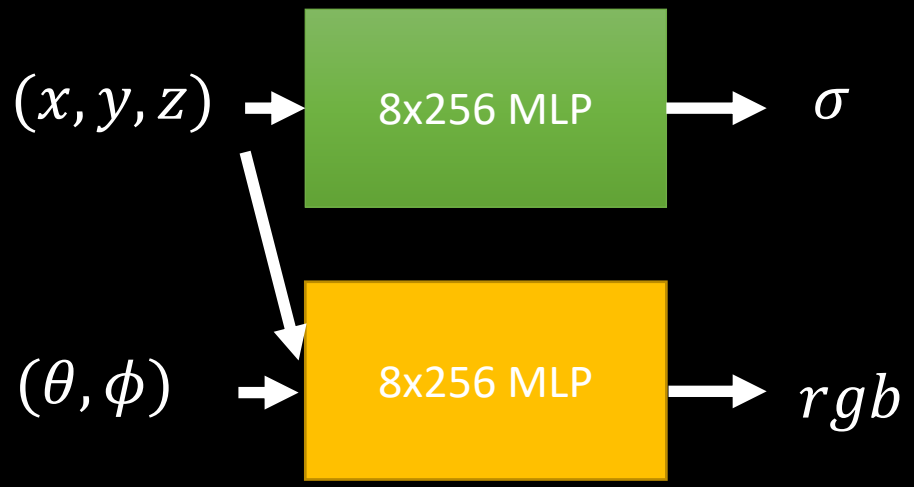


Vanilla MLP
(Remove smooth BRDF prior)

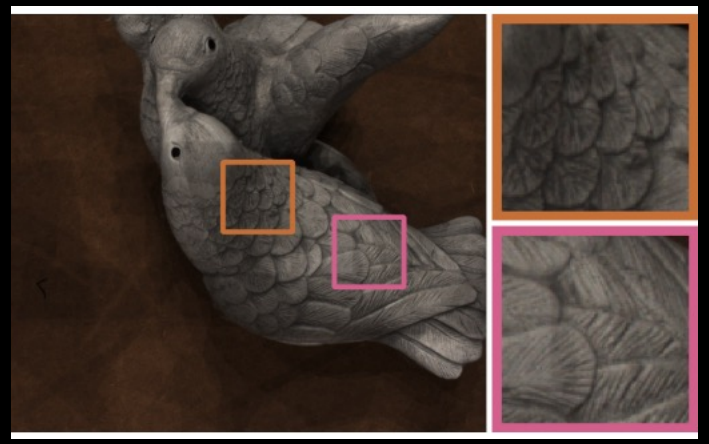

GT test view

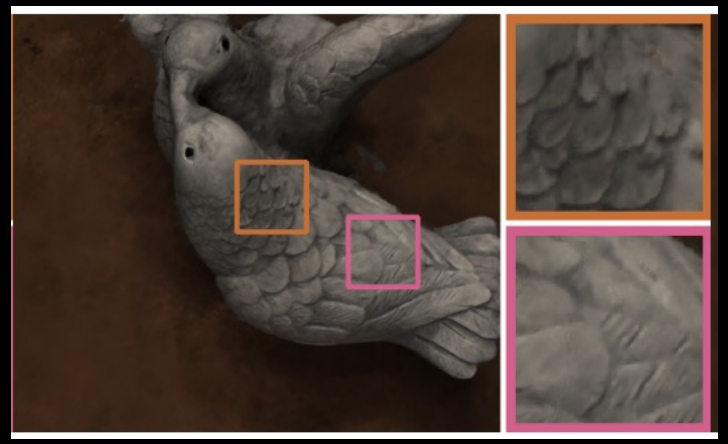

NeRF MLP

Vanilla MLP

# Resolution issue for 360 capture of large scenes

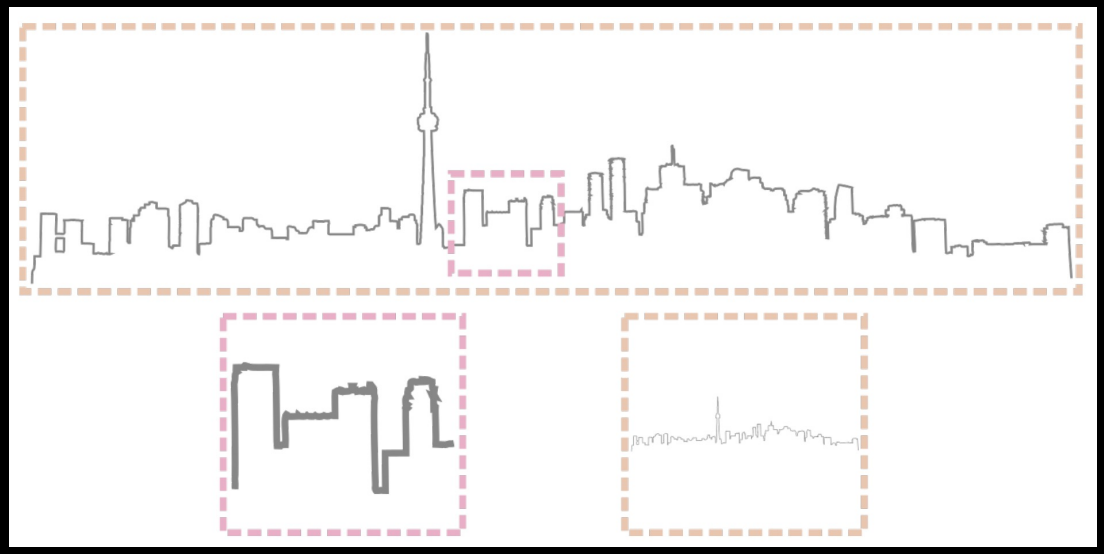

(a) bounding volume for the truck only

(b) bounding volume for the entire scene

# Our solution: inverted sphere parametrization

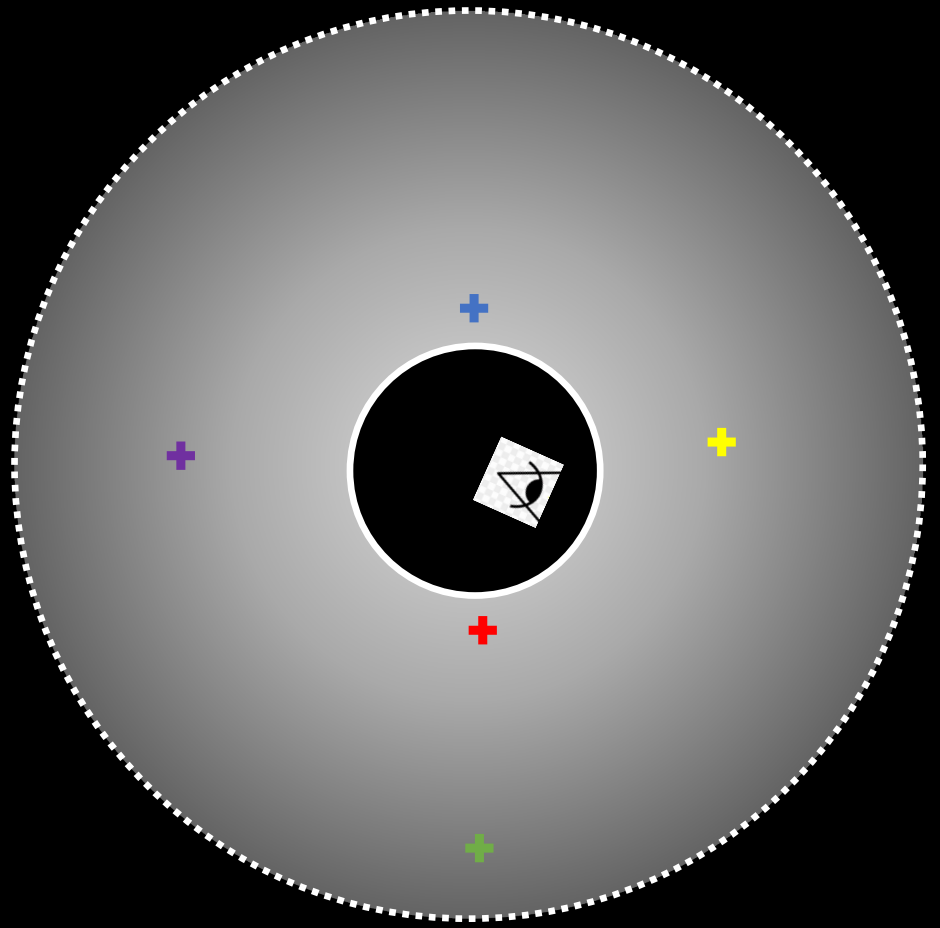

Unbounded region outside unit sphere

- Map unbounded space outside a sphere into a unit cube.
- Farther-away points get more “squeezed” after mapping.

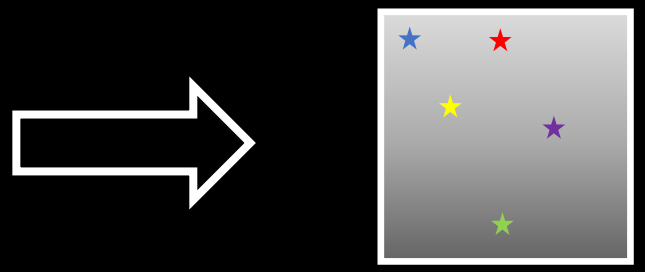

Bounded unit cube

# Our solution: separate foreground/background

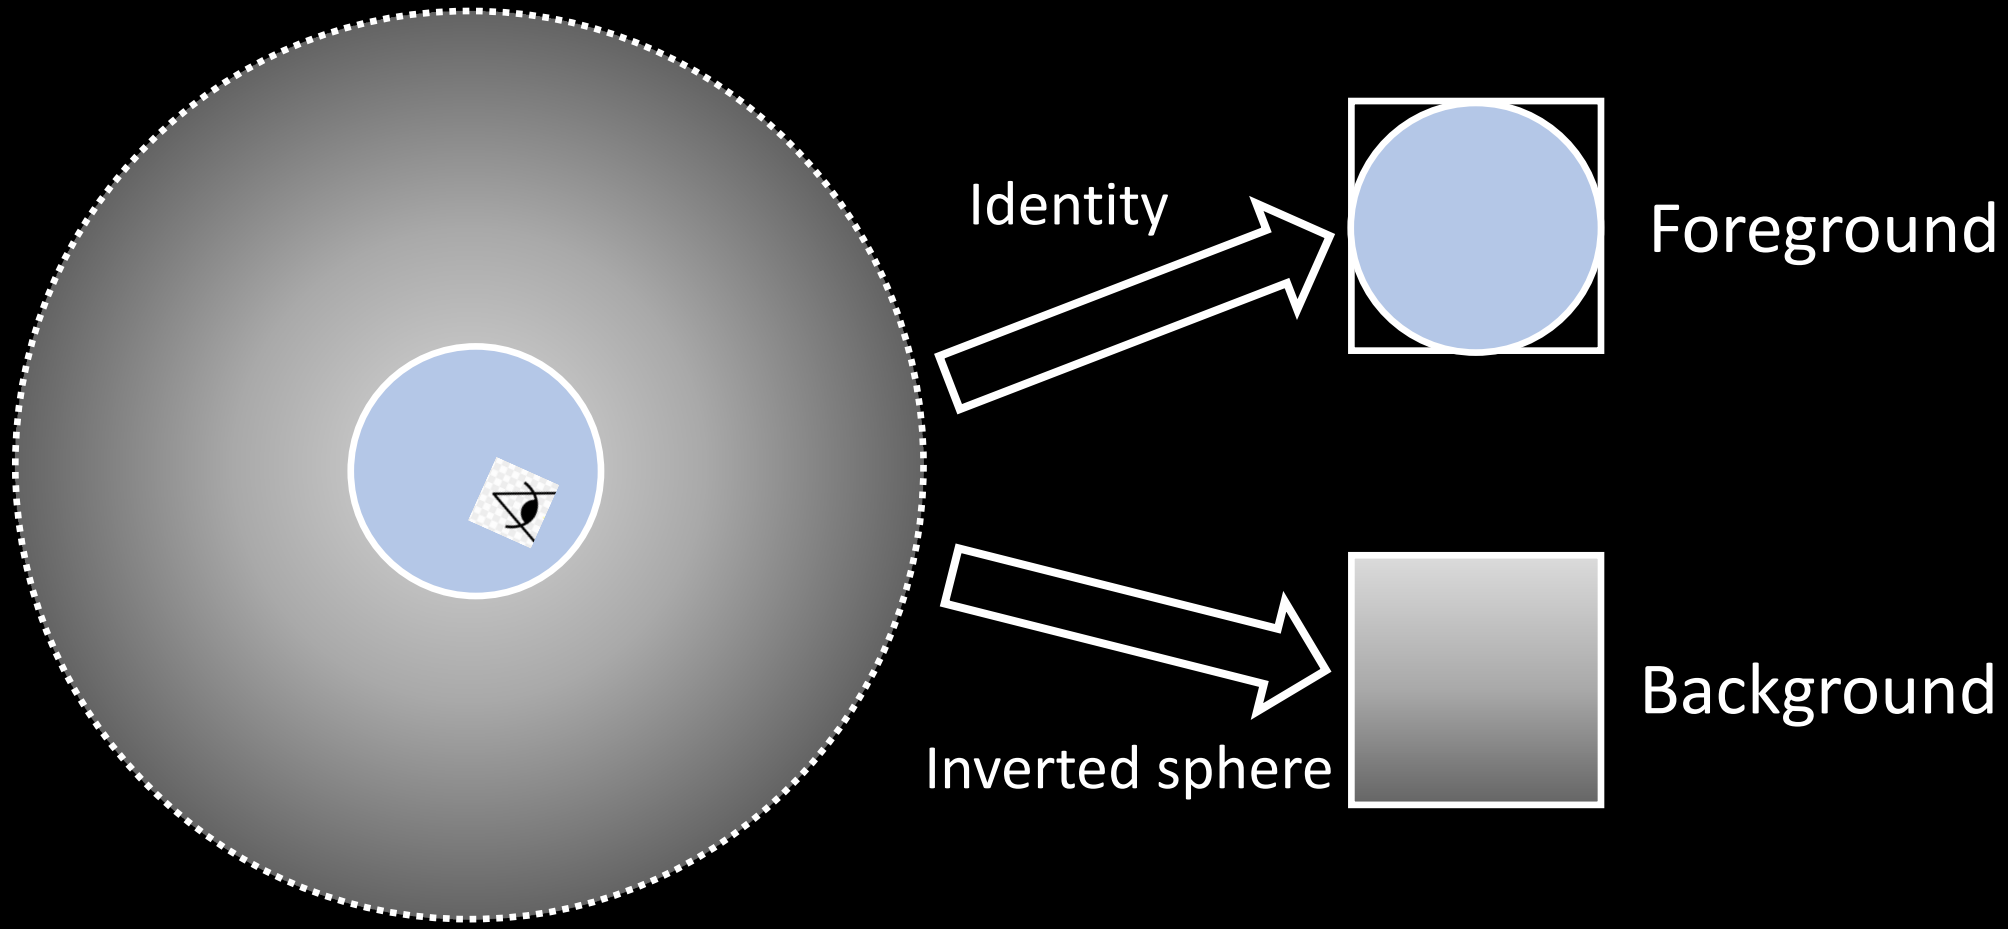


Unbounded region outside unit sphere

Bounded unit cubes

# Our solution: separate foreground/background

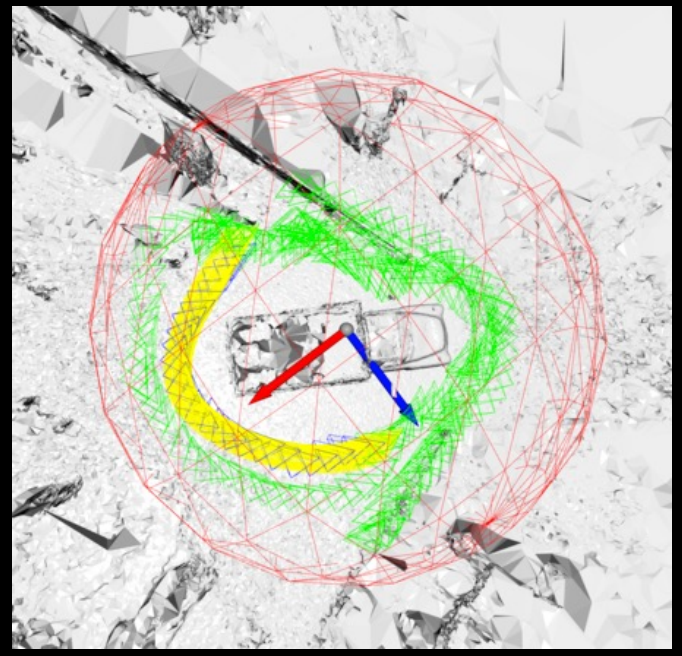

Foreground

Background

Combined

# Our results: sharper image details

NeRF

Ours

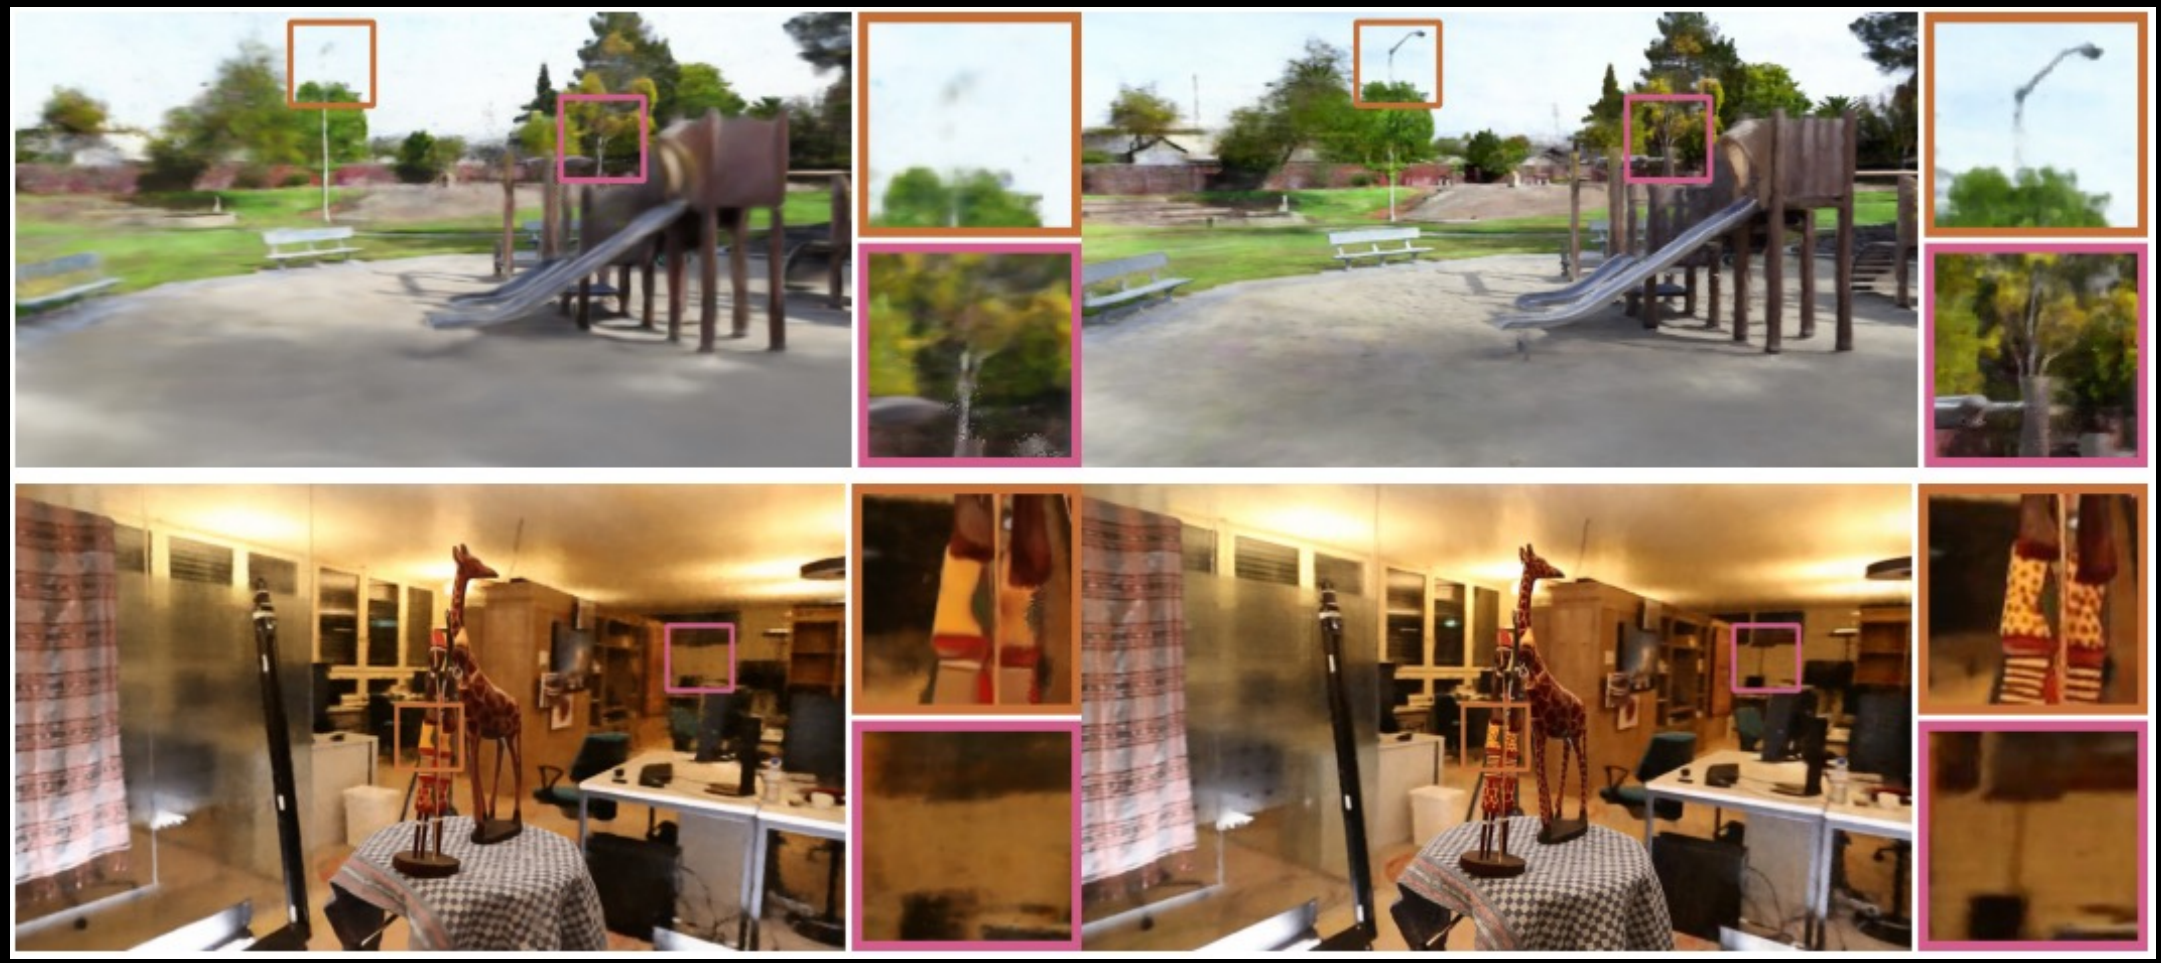

# Our results

# Take-home messages

- Shape-radiance ambiguity
  - Easily optimizable soft shape
  - Smooth BRDF prior
- Resolution issue for 360 capture of large scenes
  - Foreground/background separate modelling
  - Inverted sphere parametrization for unbounded background

Thank you!

Q & A?

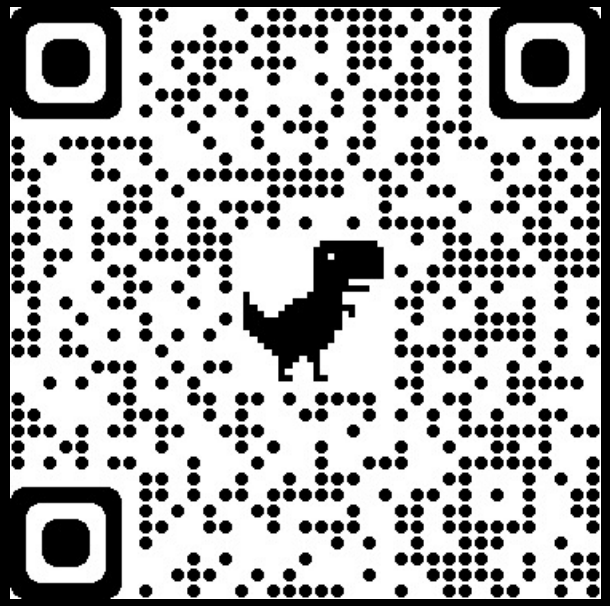

Try our code!