Neural Representations through Shadows

Kushagra Tiwary

Camera Culture, MIT Media Lab

Tiwary*, Klinghoffer*, and Raskar; Towards Learning Neural Representations from Shadows, *ECCV 2022*

Klinghoffer*, Somasundaram*, Tiwary*, and Raskar; Physics vs. Learned Priors: Rethinking Camera and Algorithm Design for Task-Specific Imaging; *ICCP 2022*

3D Reconstruction from RGB has exploded

[Source]

Differentiable Rendering has become a dominant hammer

Mitsuba 2

PHYSICALLY BASED RENDERER

PSDR-Cuda

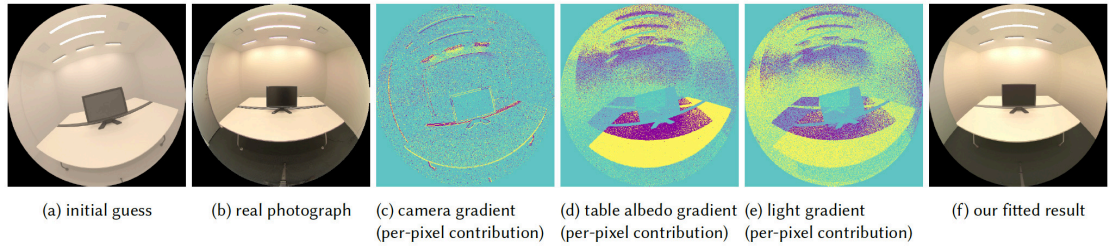

Redner

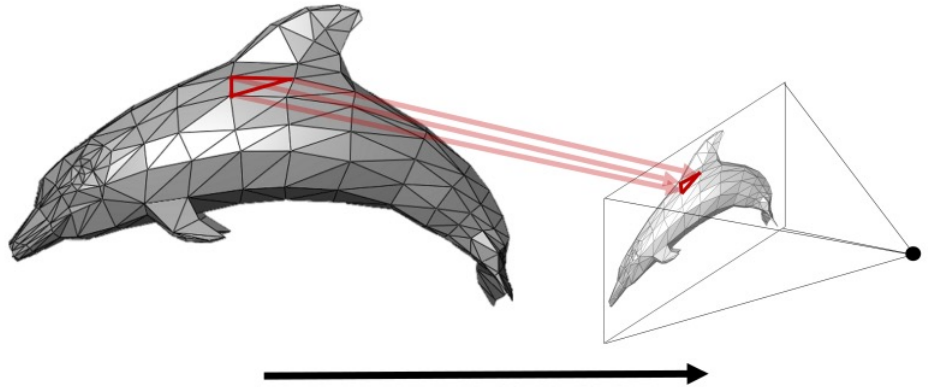

Rasterization Based

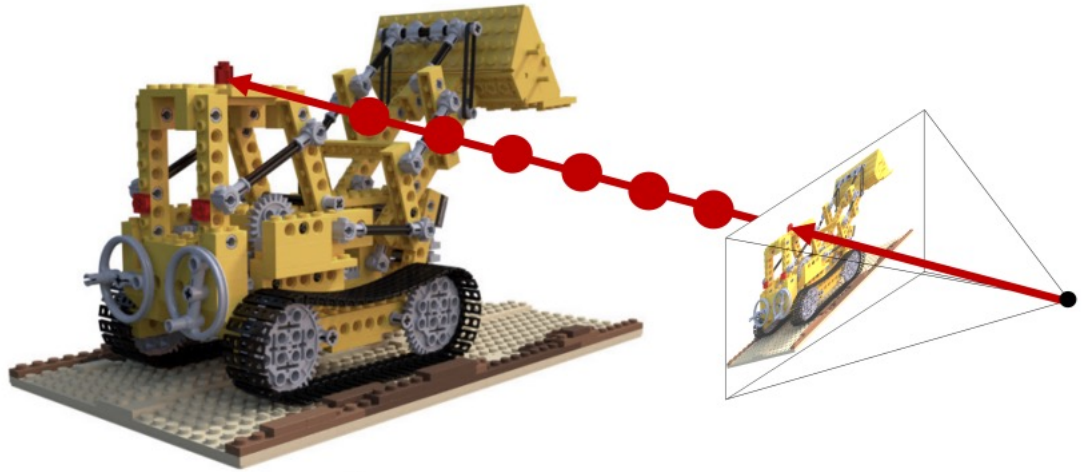

Ray Casting

Source: CVPR 2021 Tutorial

Neural Rendering has become a dominant hammer

$$pixel = \sum_i Transmittance_i \times alpha_i \times color_i$$

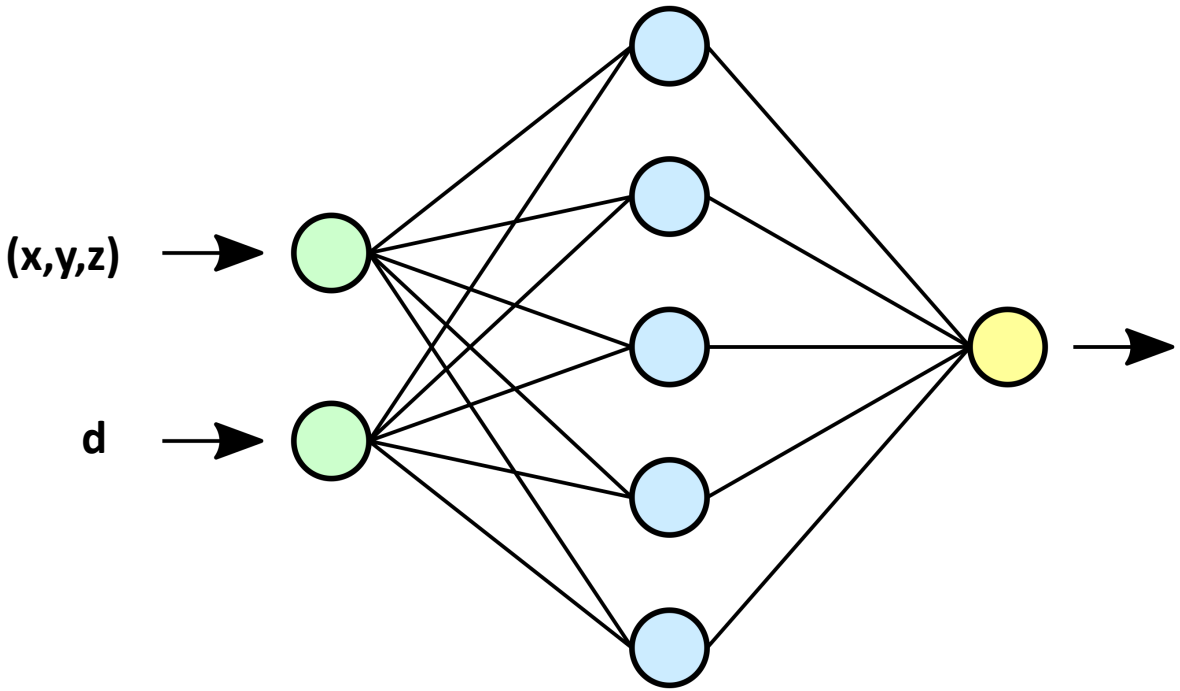

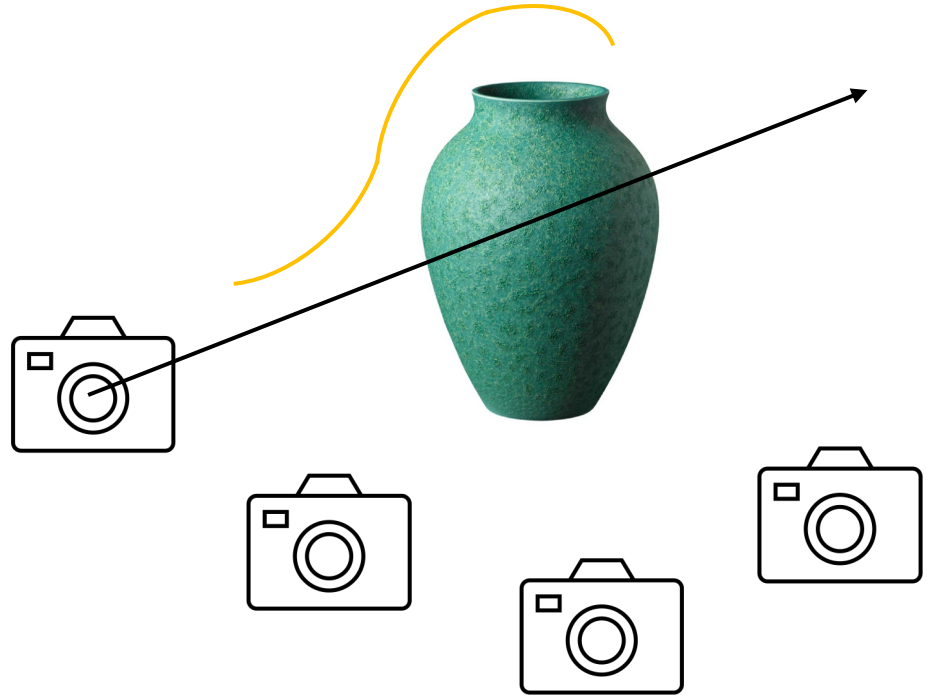

… but largely ignore physical cues present in the scene

Unconstrained Collection/Meta Learning

$$Loss(\hat{I} - I_{gt})$$

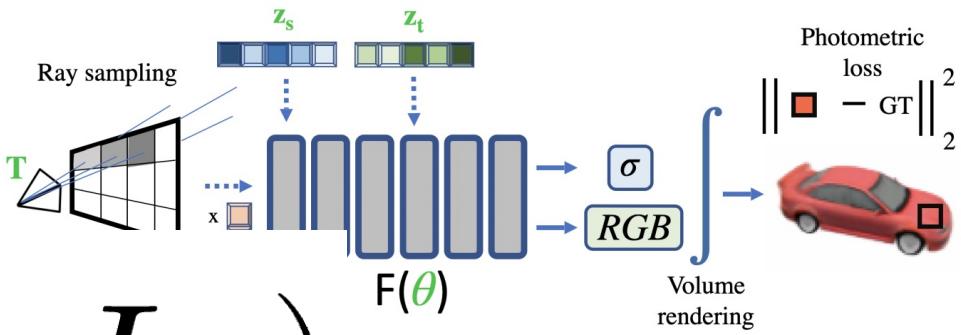

Object Priors

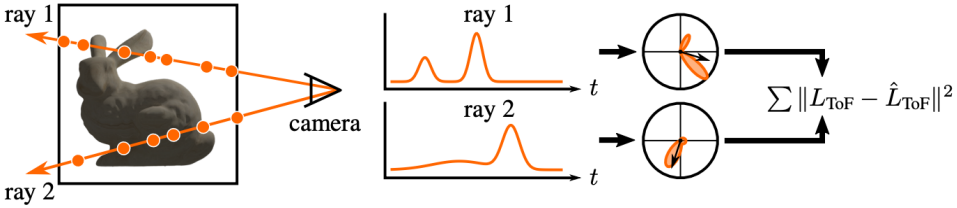

ToF + Neural Rendering

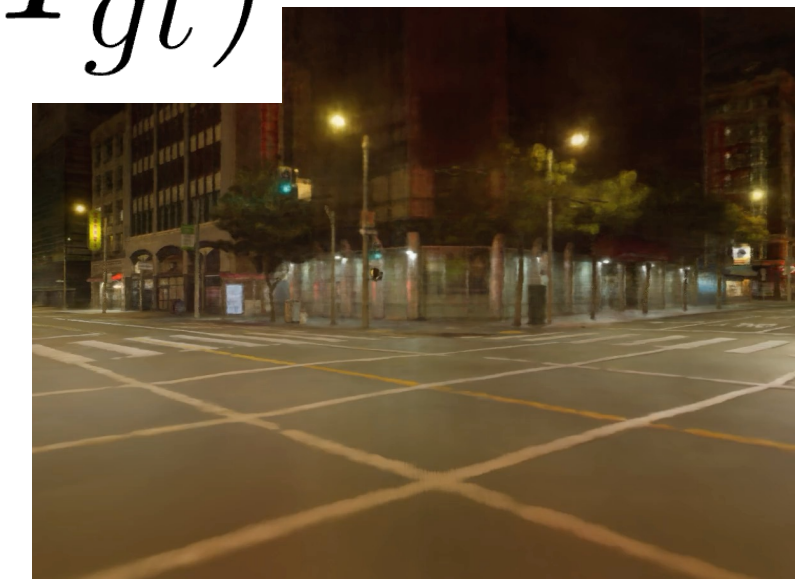

Appearance Embeddings

Over-reliance on Photometric Consistency

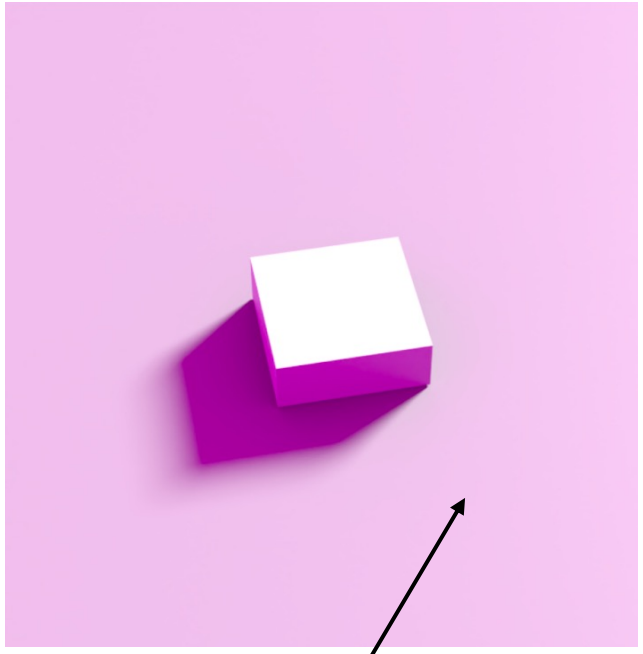

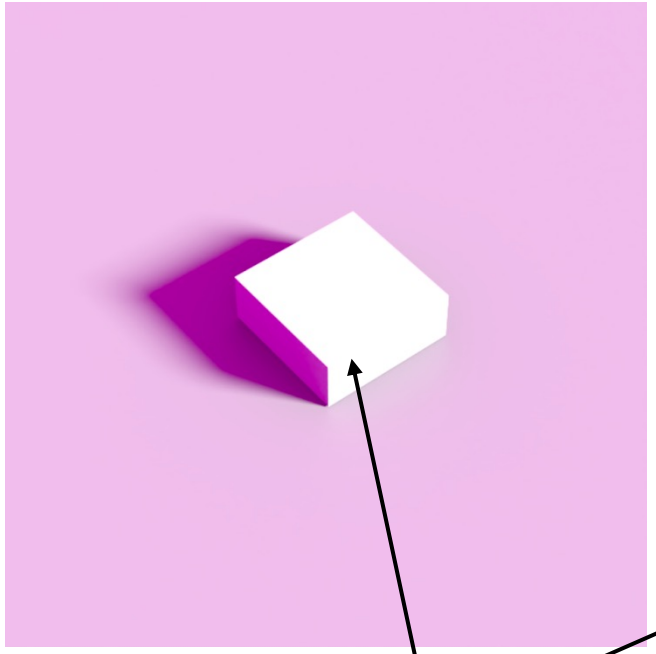

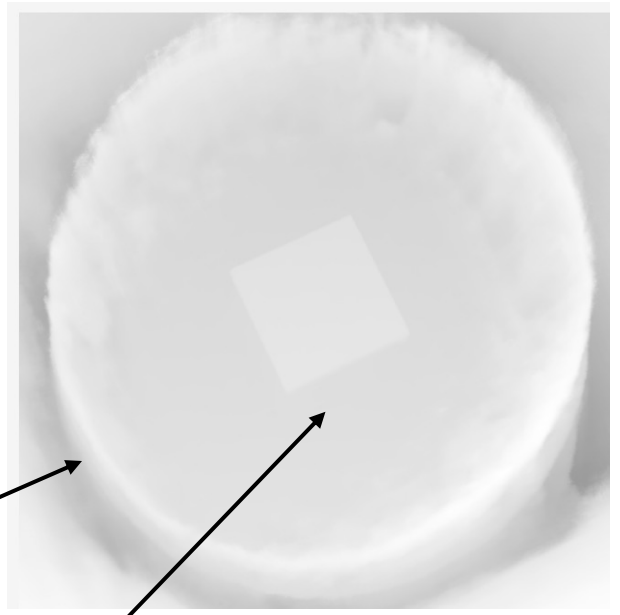

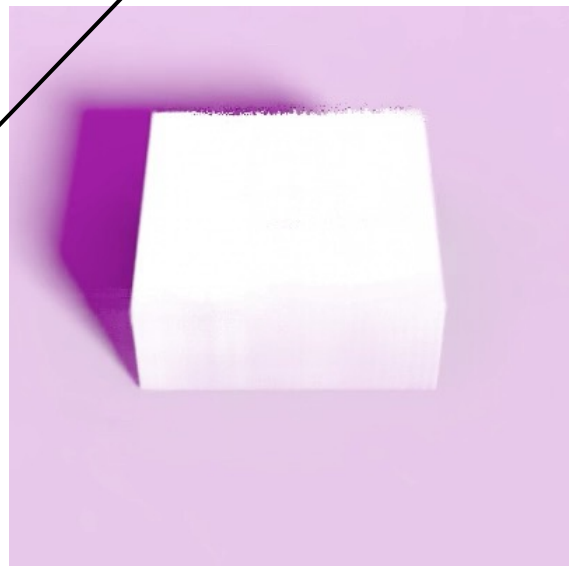

Texture Less background leads to “floaters”, hallucination or inconsistent depth maps

Texture Less Surface leads to inconsistent meshes

Over-reliance on Photometric Consistency

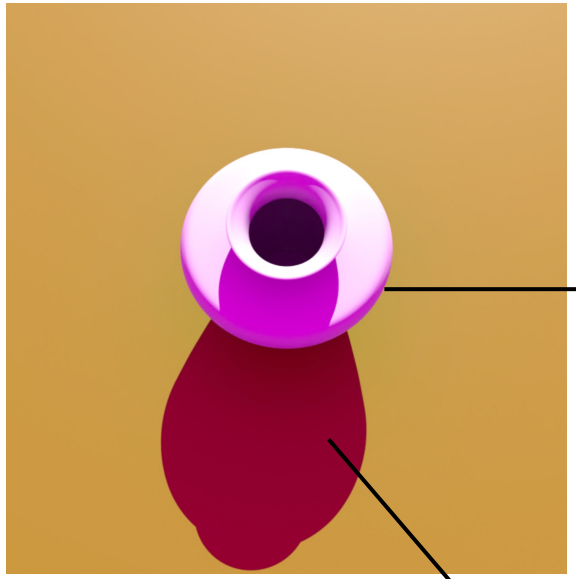

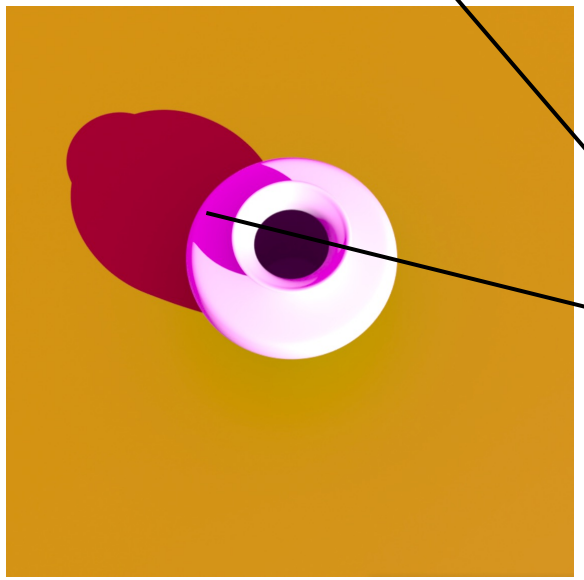

Poorly Sampled Vertical Faces exposes heavy reliance on photometric cues

Oblique Lighting hides geometry in shadows **ignored** by SOTA NeRF models

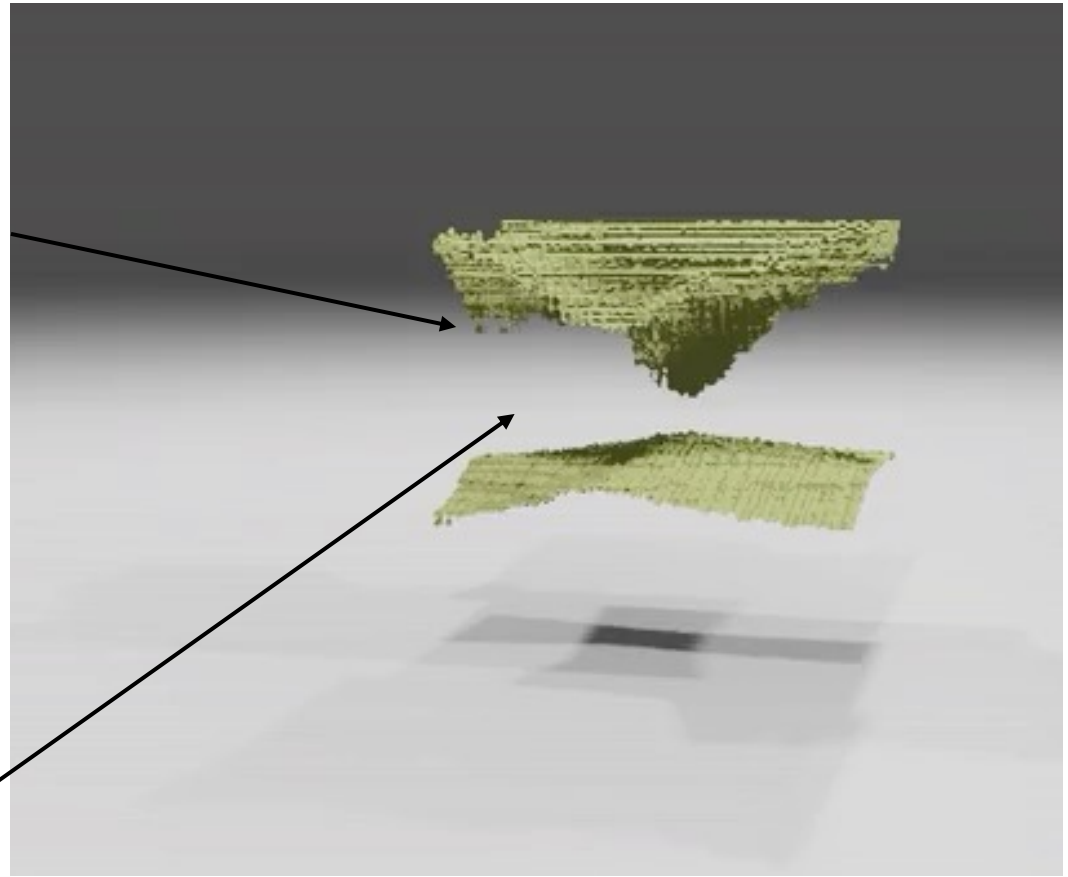

Second Order Cues Provide Vital Information

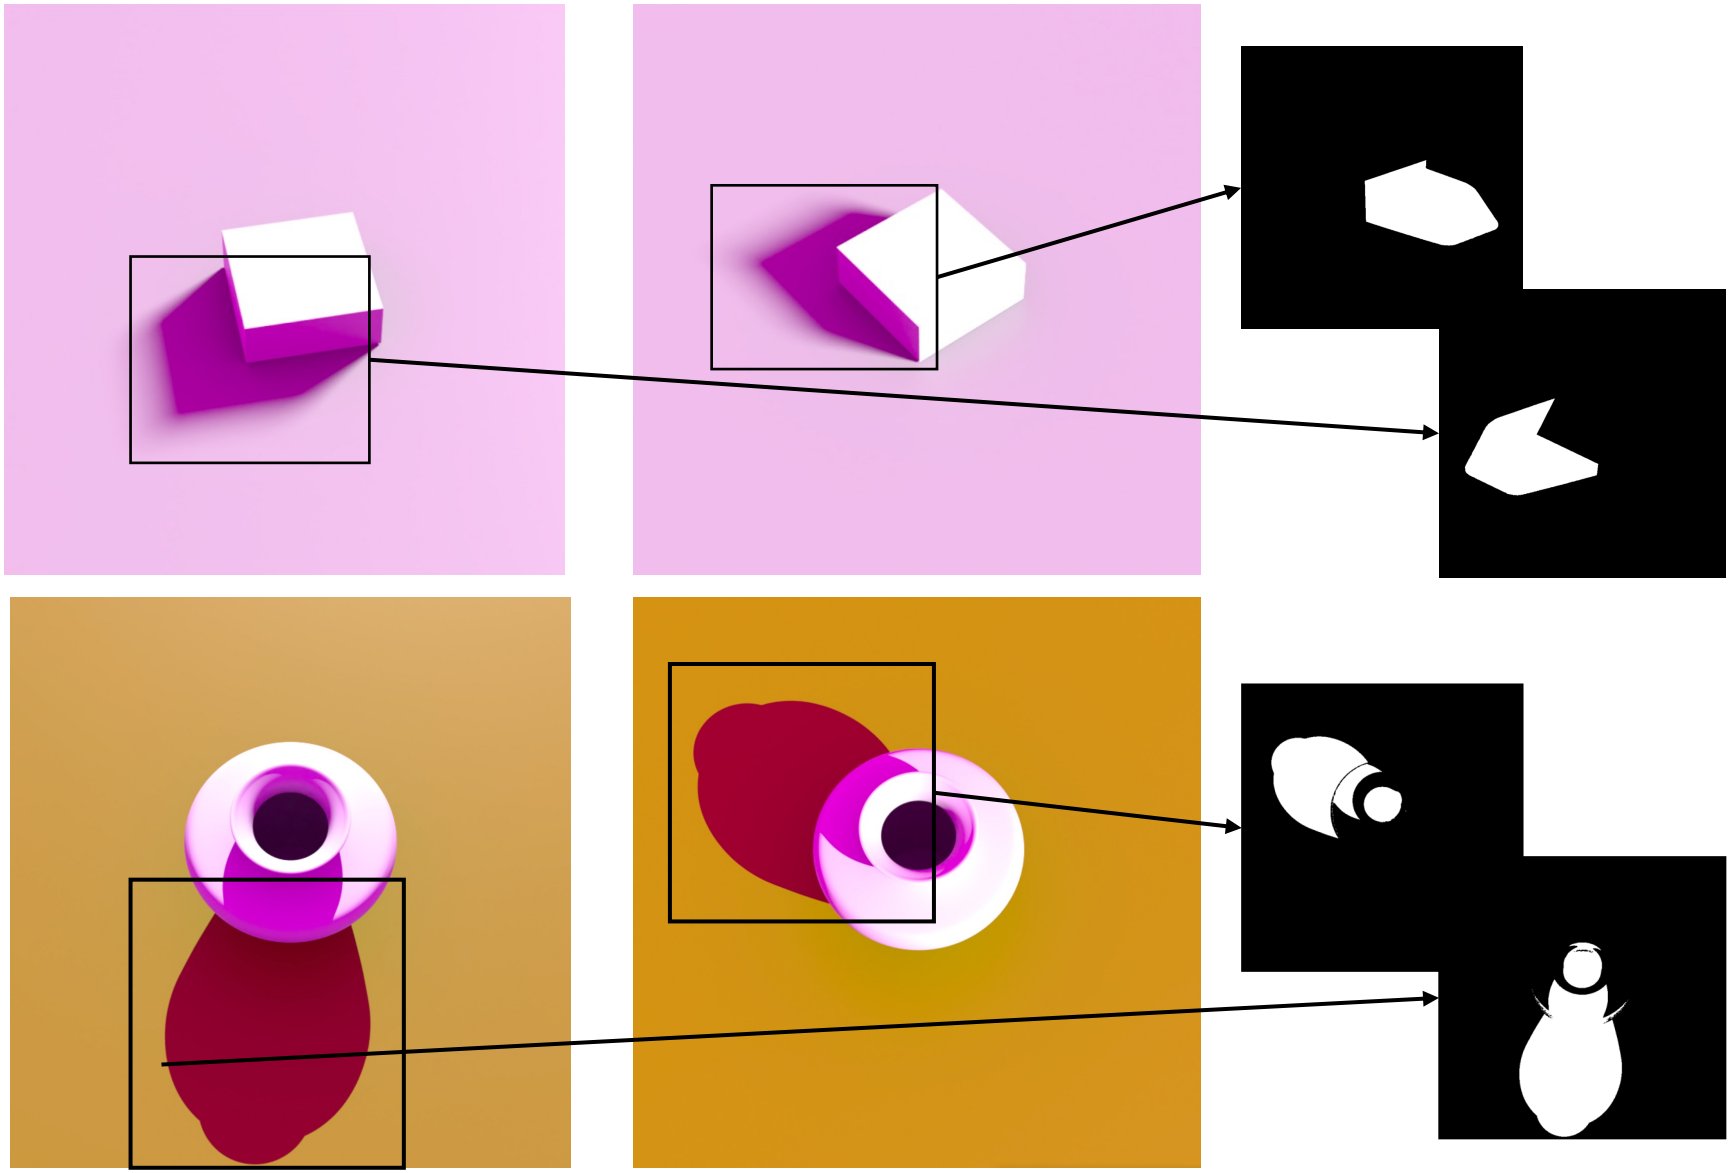

Shadows Provide a lot of Information:

- Object Concavities
- Cast Shadows → Geometry
- Independent of textures and surface reflectance models

Instead, let's exploit Second Order Cues- shadows

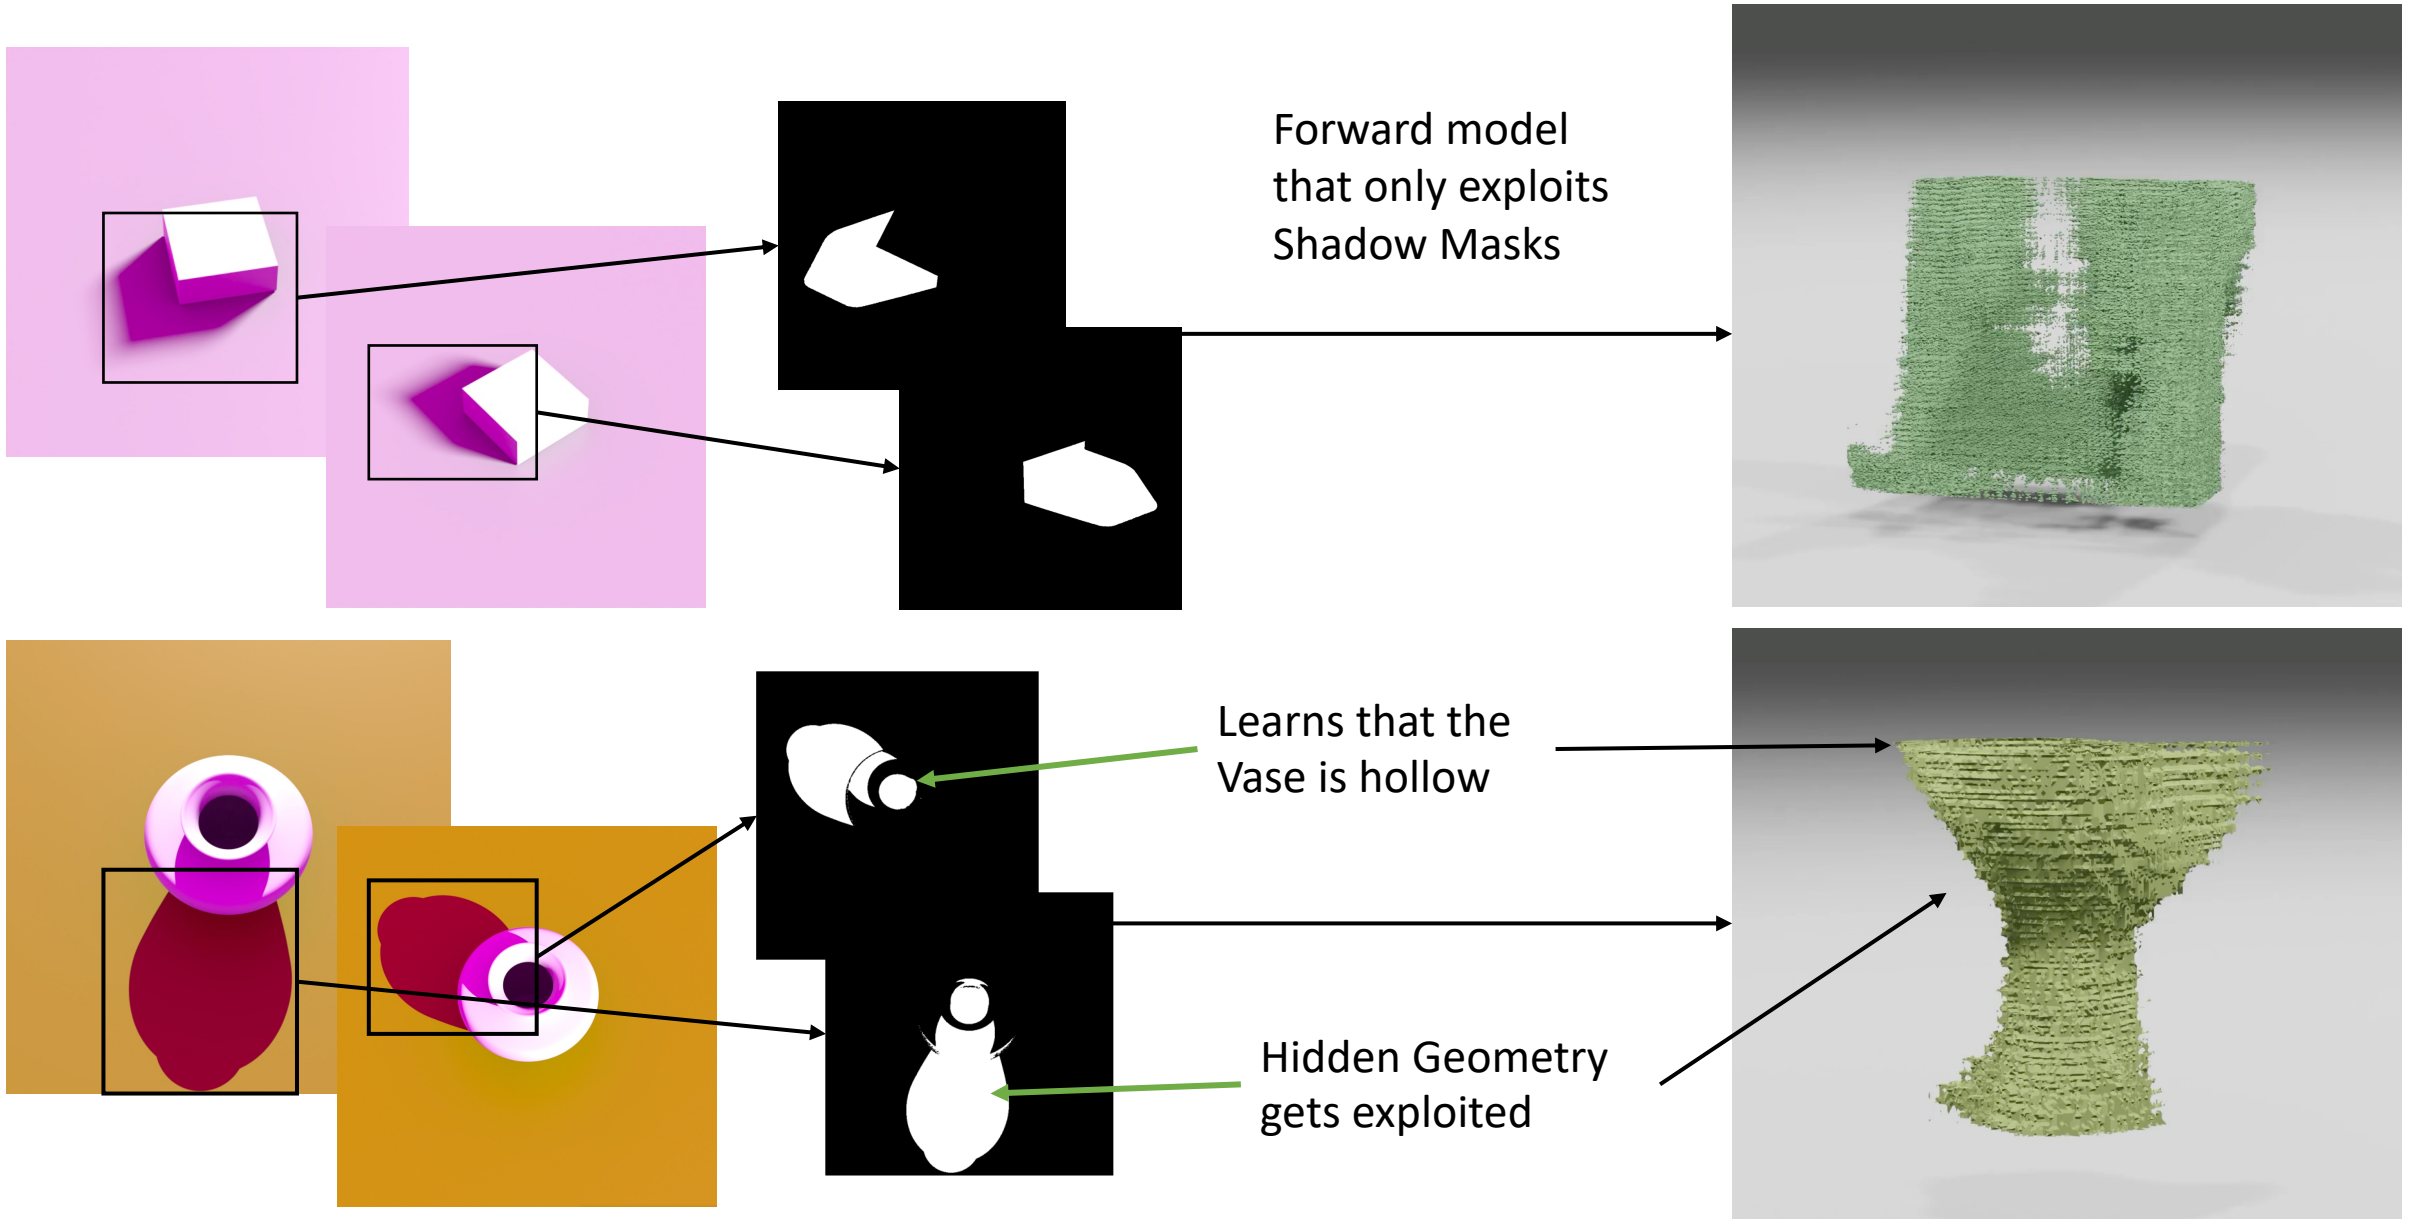

Overall Approach to Exploit only Shadows

Binary Shadow Masks captured with varying camera position and fixed lighting

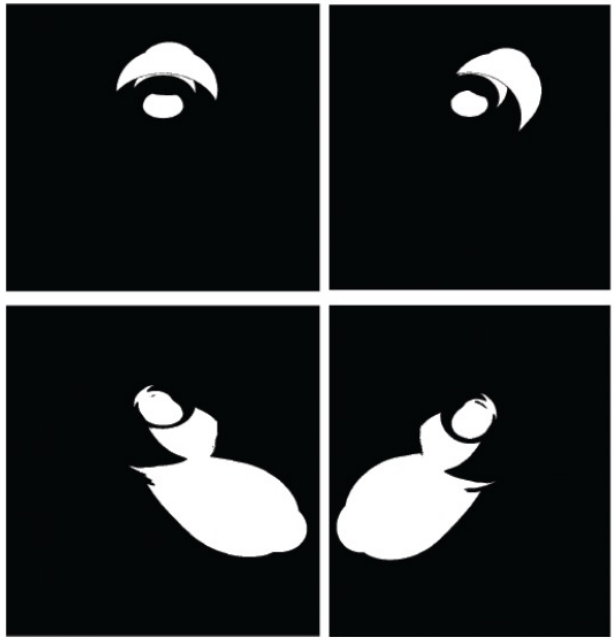

Proposed approach to exploit shadow cues in the scene

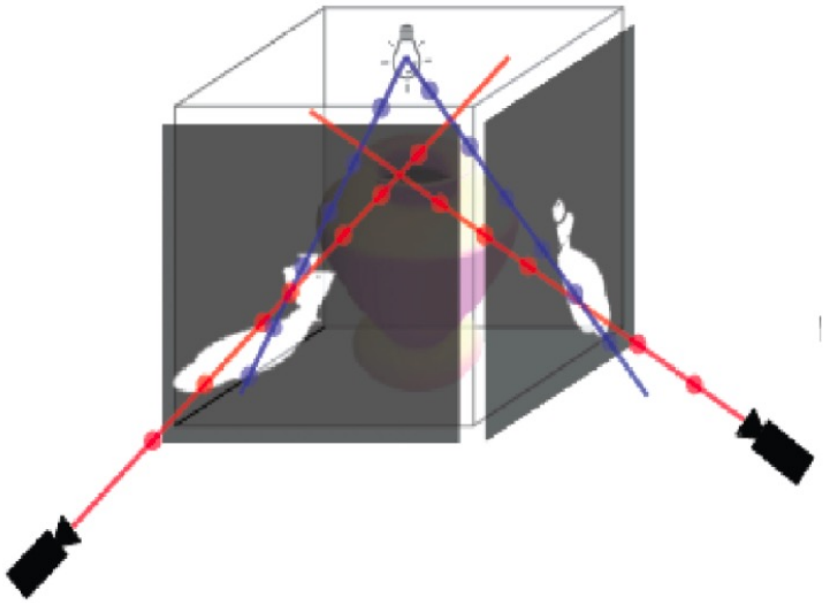

Estimated Depth, Shadow Mask, Disparity Map, and Mesh only through binary shadows

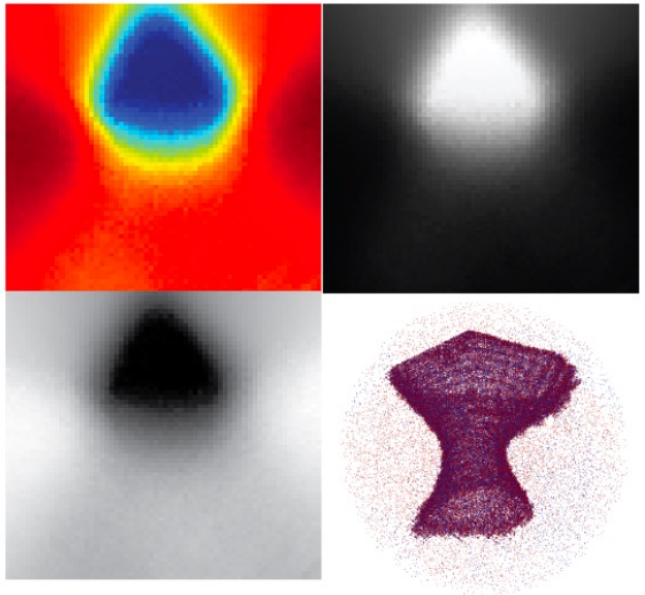

What are Shadows?

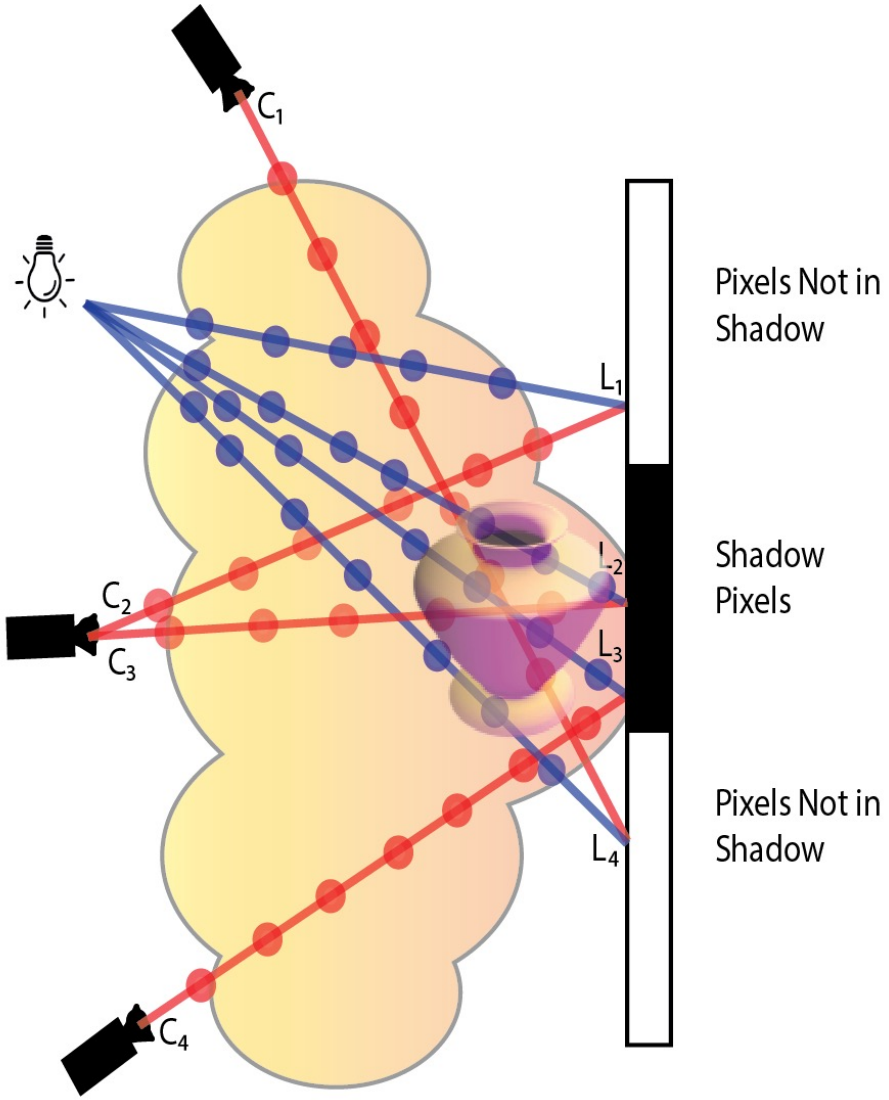

All points in the world without a direct path to the *light source* are defined to be in **shadow**.

Quick Primer on Shadow Mapping

Shadow Map: Distance to the scene from the light's perspective

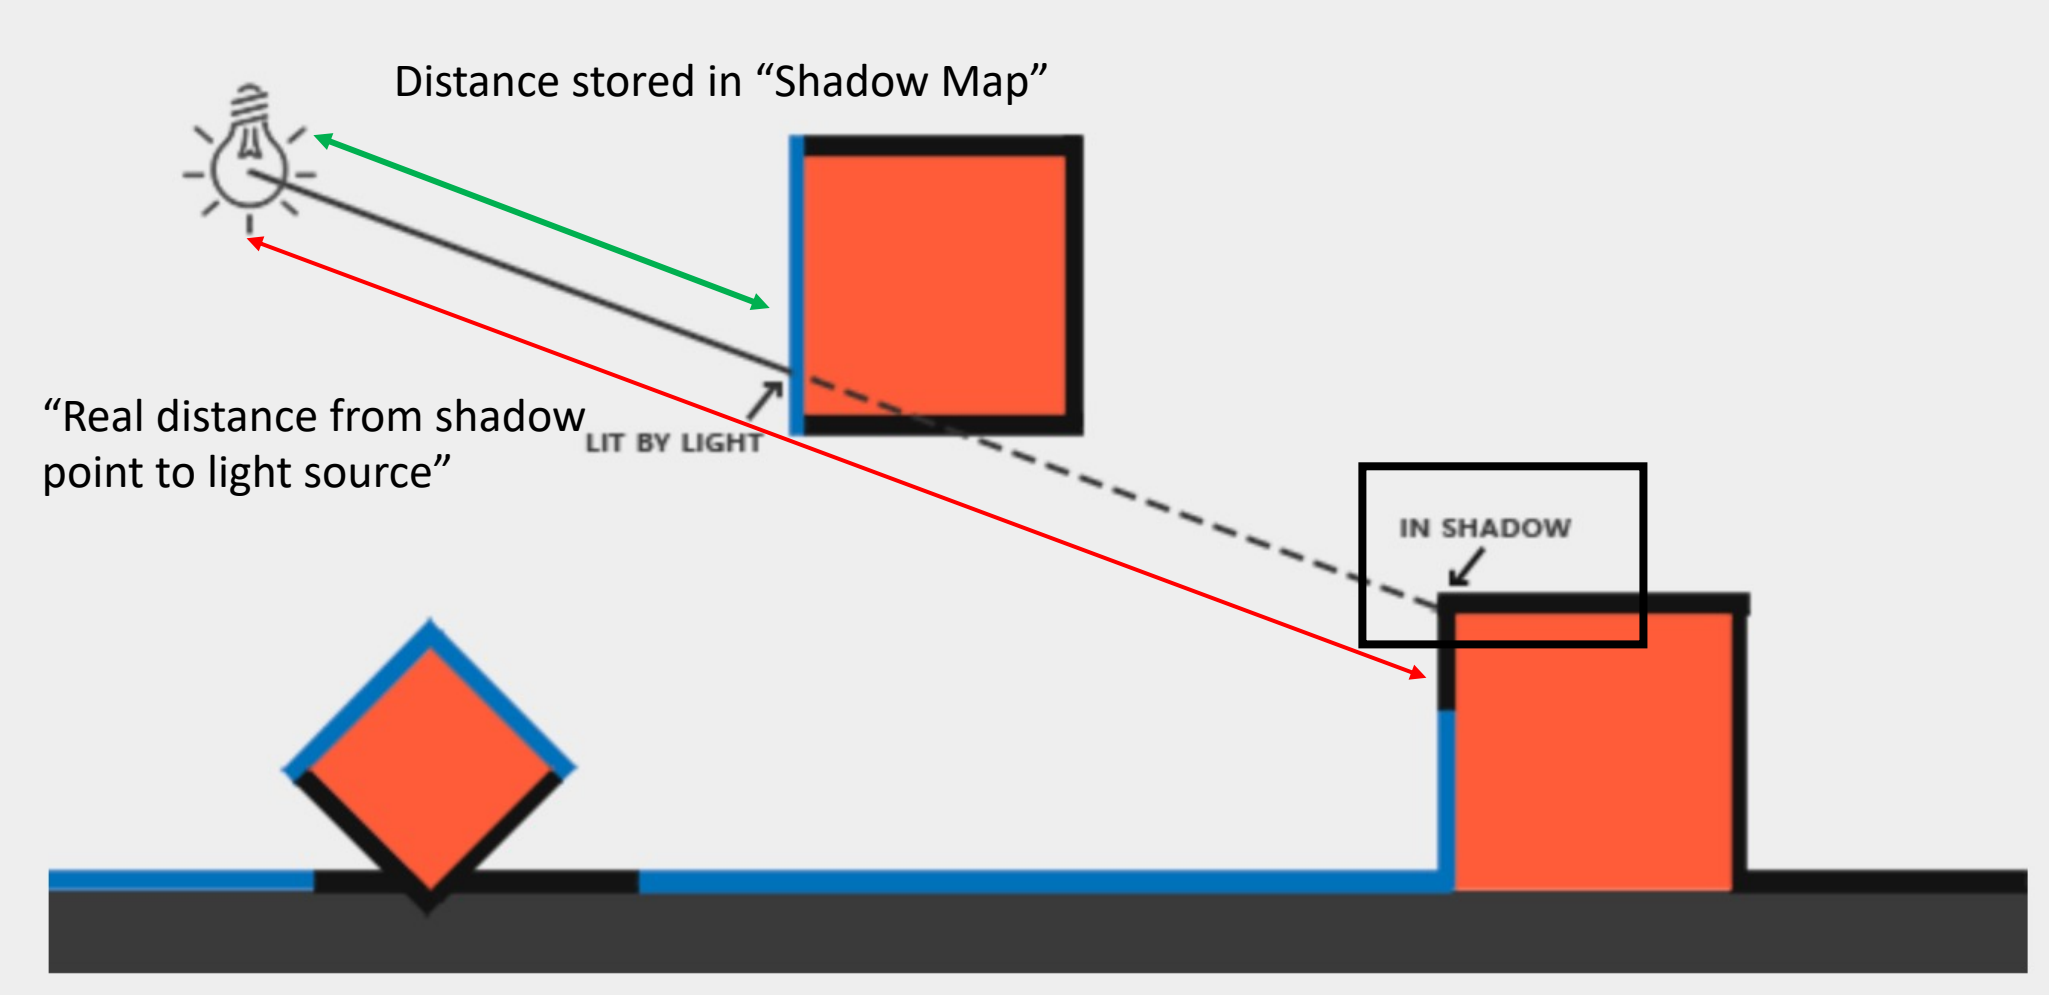

Source

Shadow Mapping

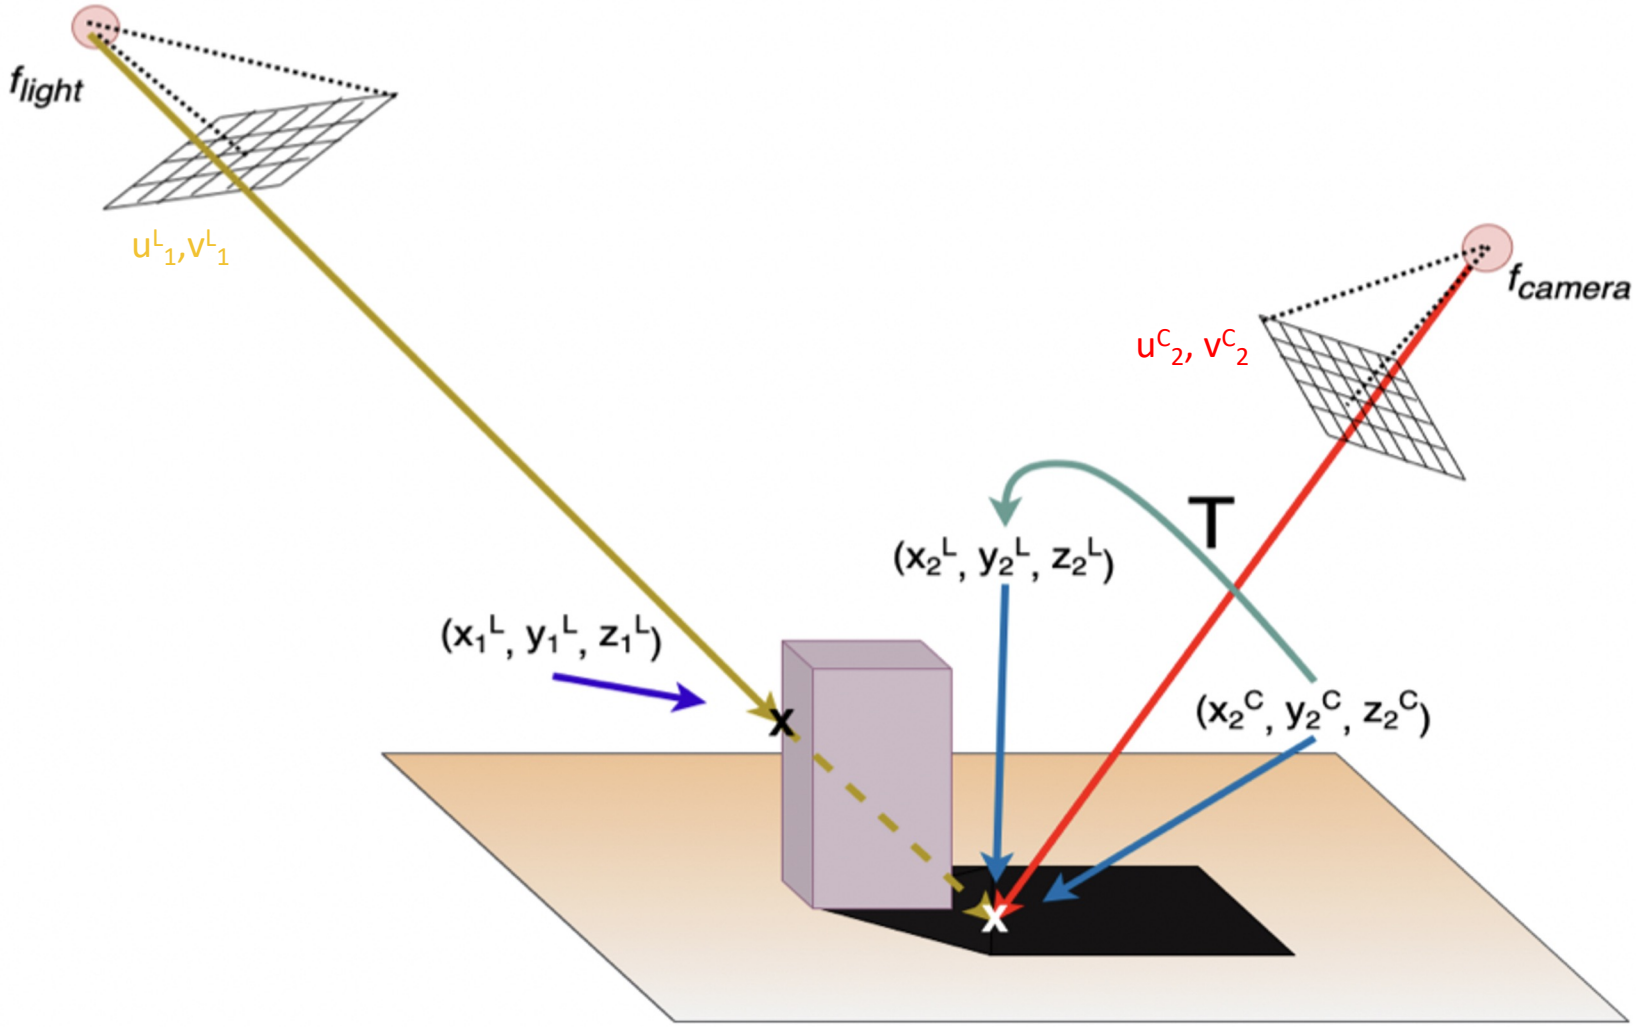

Let's consider pixels:

- $(u^C_2, v^C_2, 1) \rightarrow (x^C_2, y^C_2, z^C_2)$
- $(u^L_1, v^L_1, 1) \rightarrow (x^L_2, y^L_2, z^L_2)$
- *Function **F**: pixel -> Depth at Pixel*
- *Transformation **T: from_camera_to_light***

Shadow Mapping:

1. $F_{camera}((u^C_2, v^C_2, 1)) = (x^C_2, y^C_2, z^C_2)$
2. $F_{light}((u^L_1, v^L_1, 1)) = (x^L_2, y^L_2, z^L_2)$
3. $T(x^C_2, y^C_2, z^C_2) = (x^L_2, y^L_2, z^L_2)$
4. If $z^L_1 < z^L_2$ then point (x^C_2, y^C_2, z^C_2) is IN shadow.

Differentiable Shadows Forward Model

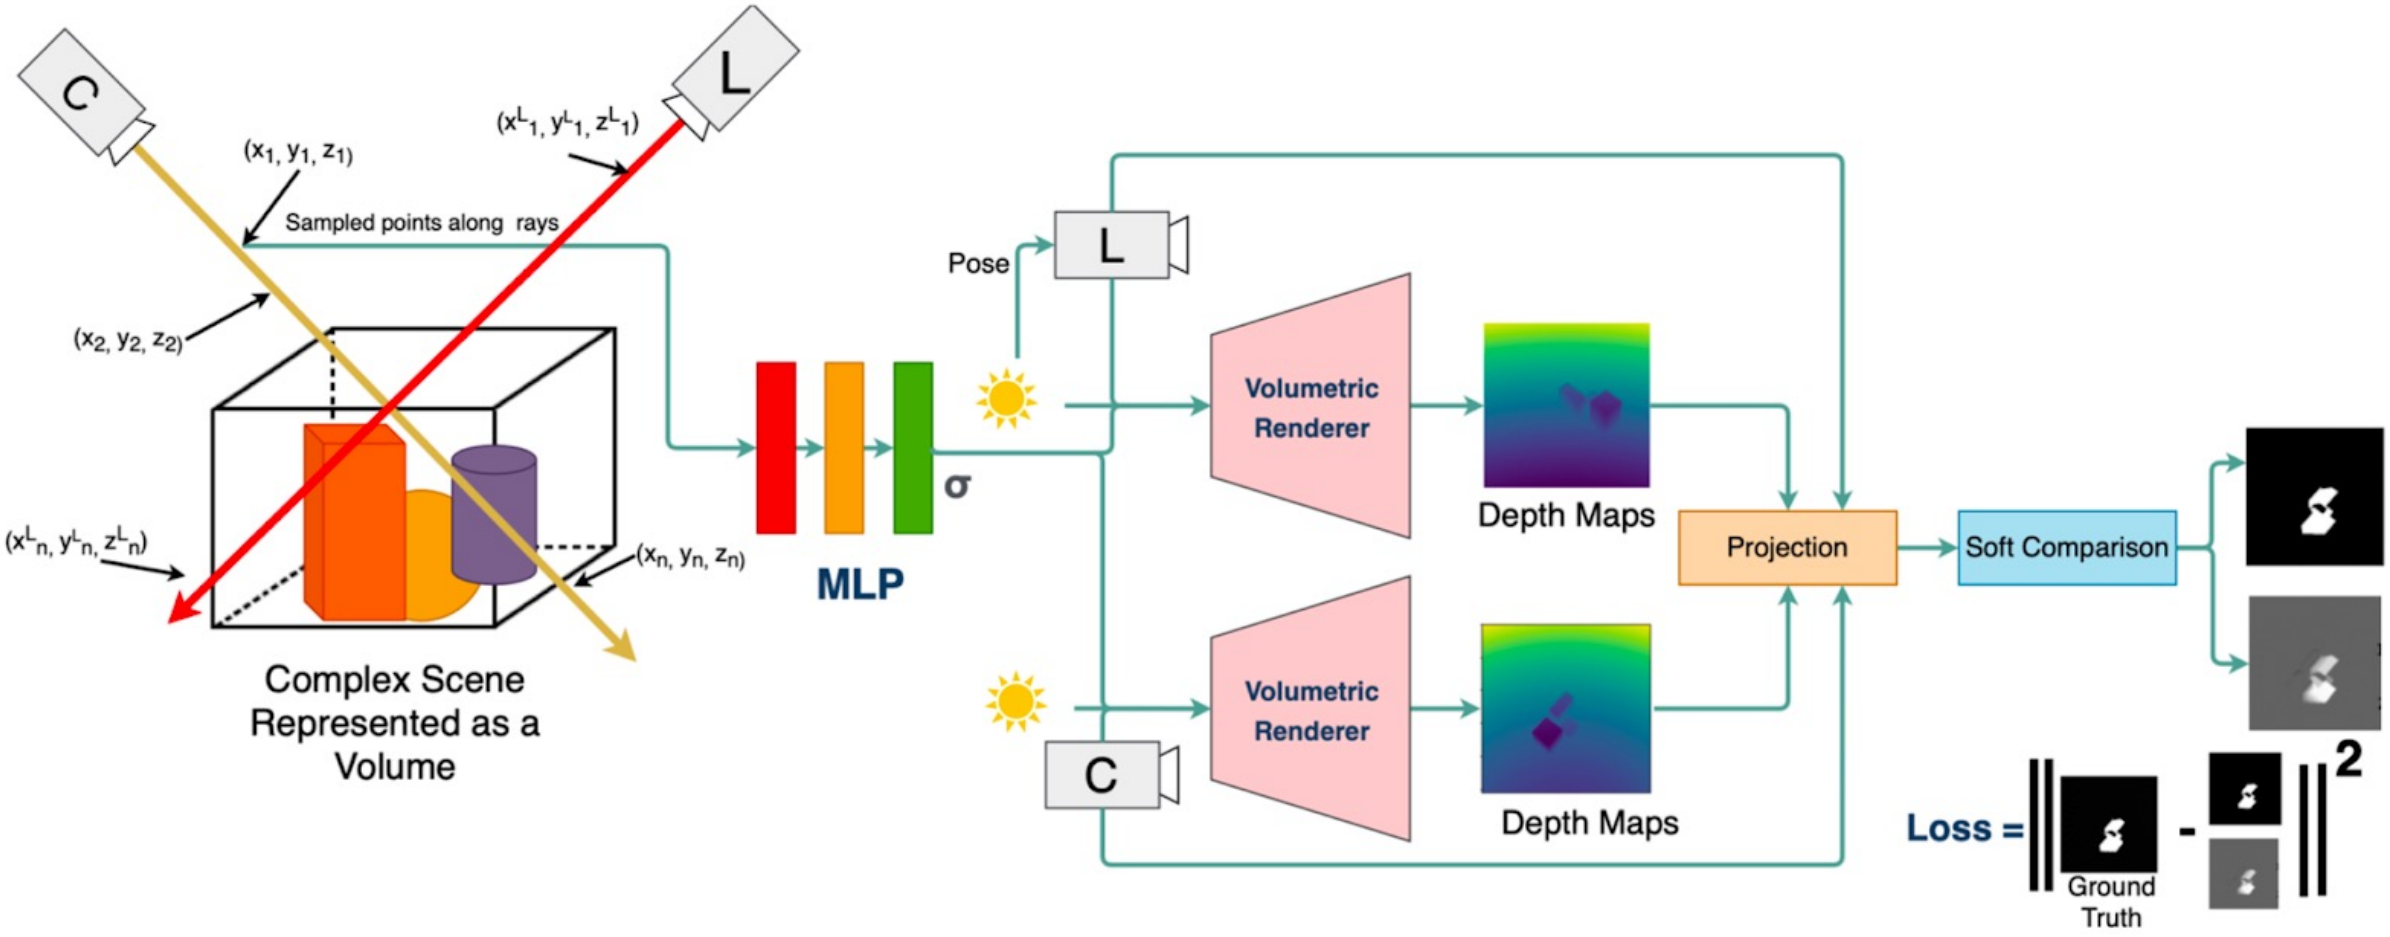

Evolution of Depth and Shadow Masks on Validation Data

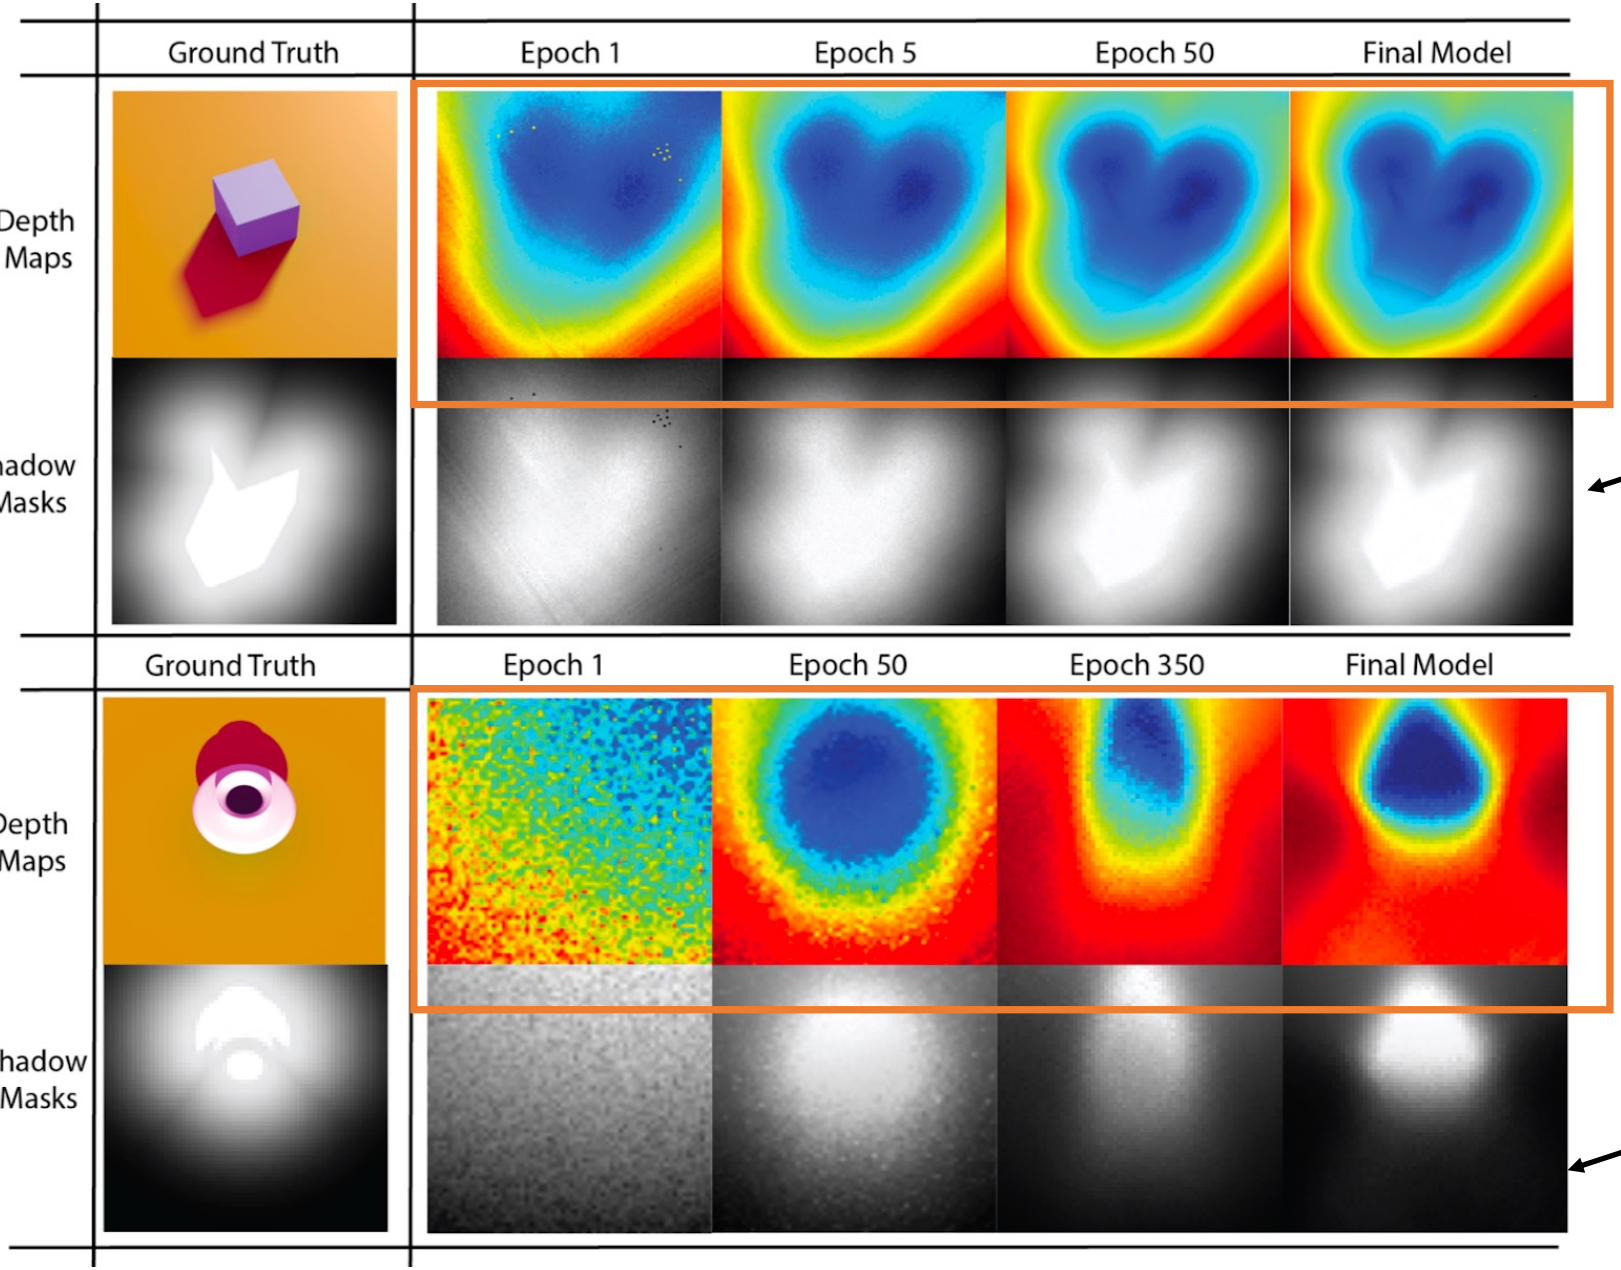

Loss is minimizing the error between Shadow Masks

Loss is minimizing the error between Shadow Masks

Recovering Poorly Sampled Vertical Surfaces with Neural Shadow Fields

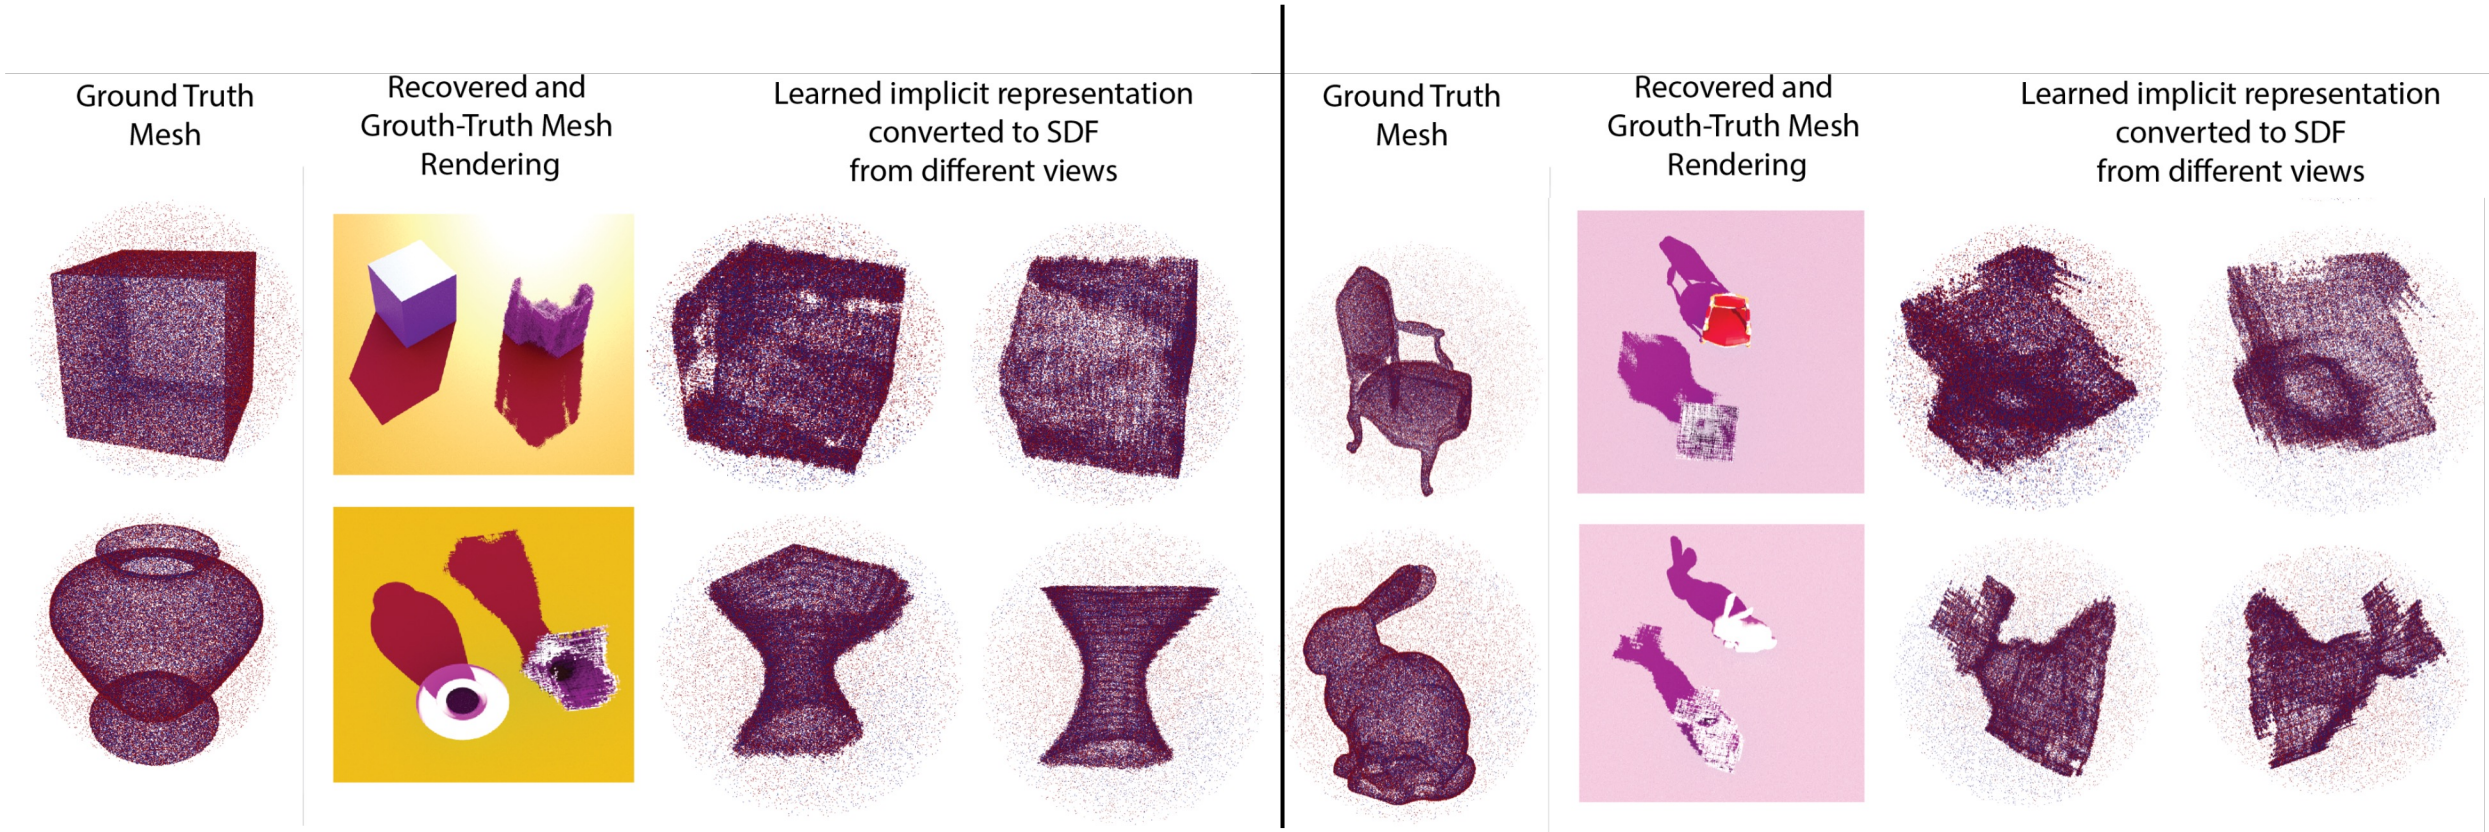

Scene	RMSE Shadow Mesh	RMSE Vanilla NeRF
Cuboid	**0.0078**	0.097
Vase	**0.010**	0.0.011
Bunny	0.0109	**0.0106**
Chair	**0.0092**	0.0096

Real World Scene: Hand Imprint

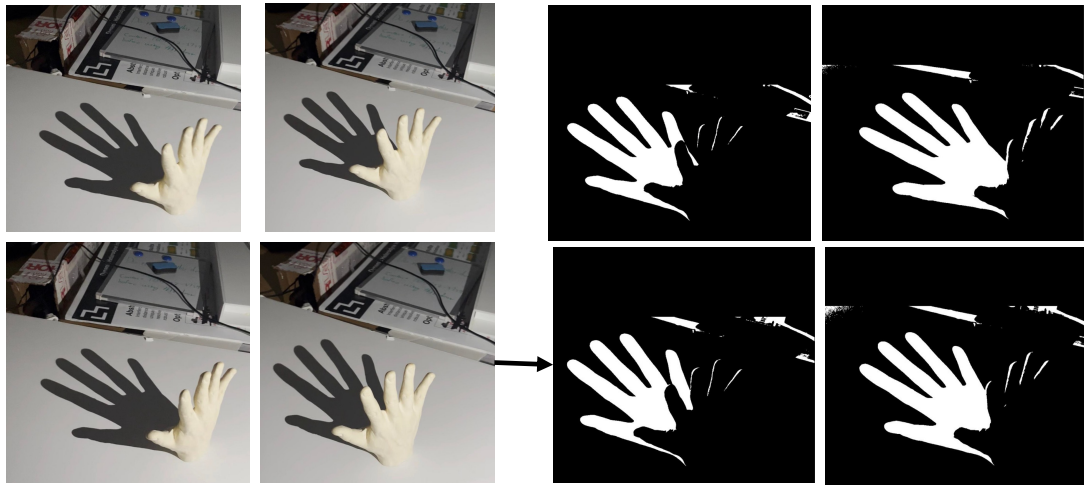

Camera View

Threshold Shadows

Light View

Light Shadow forced to be 0

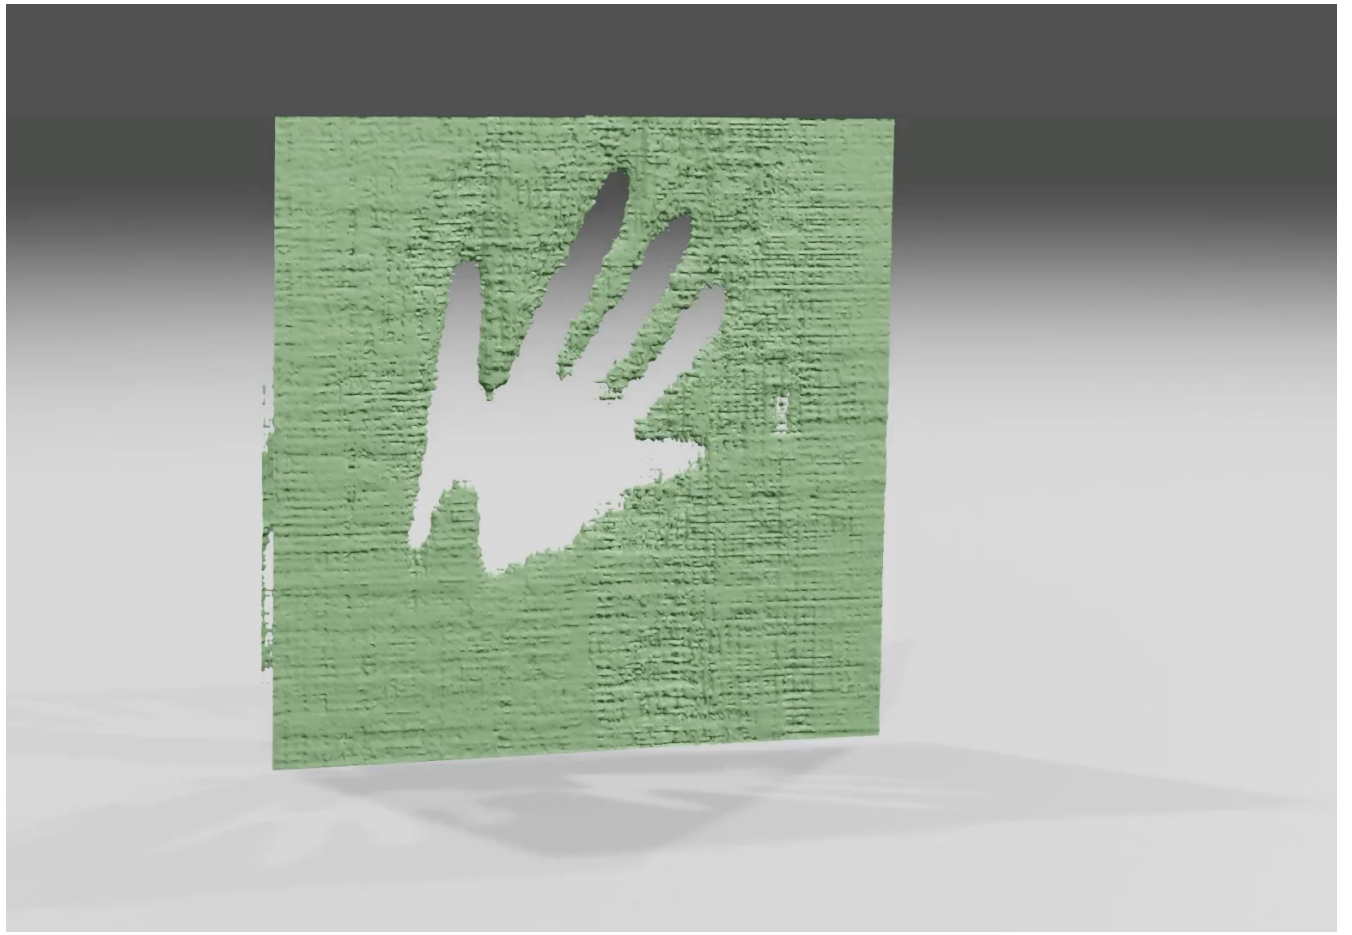

Real World Scene: Hand Imprint

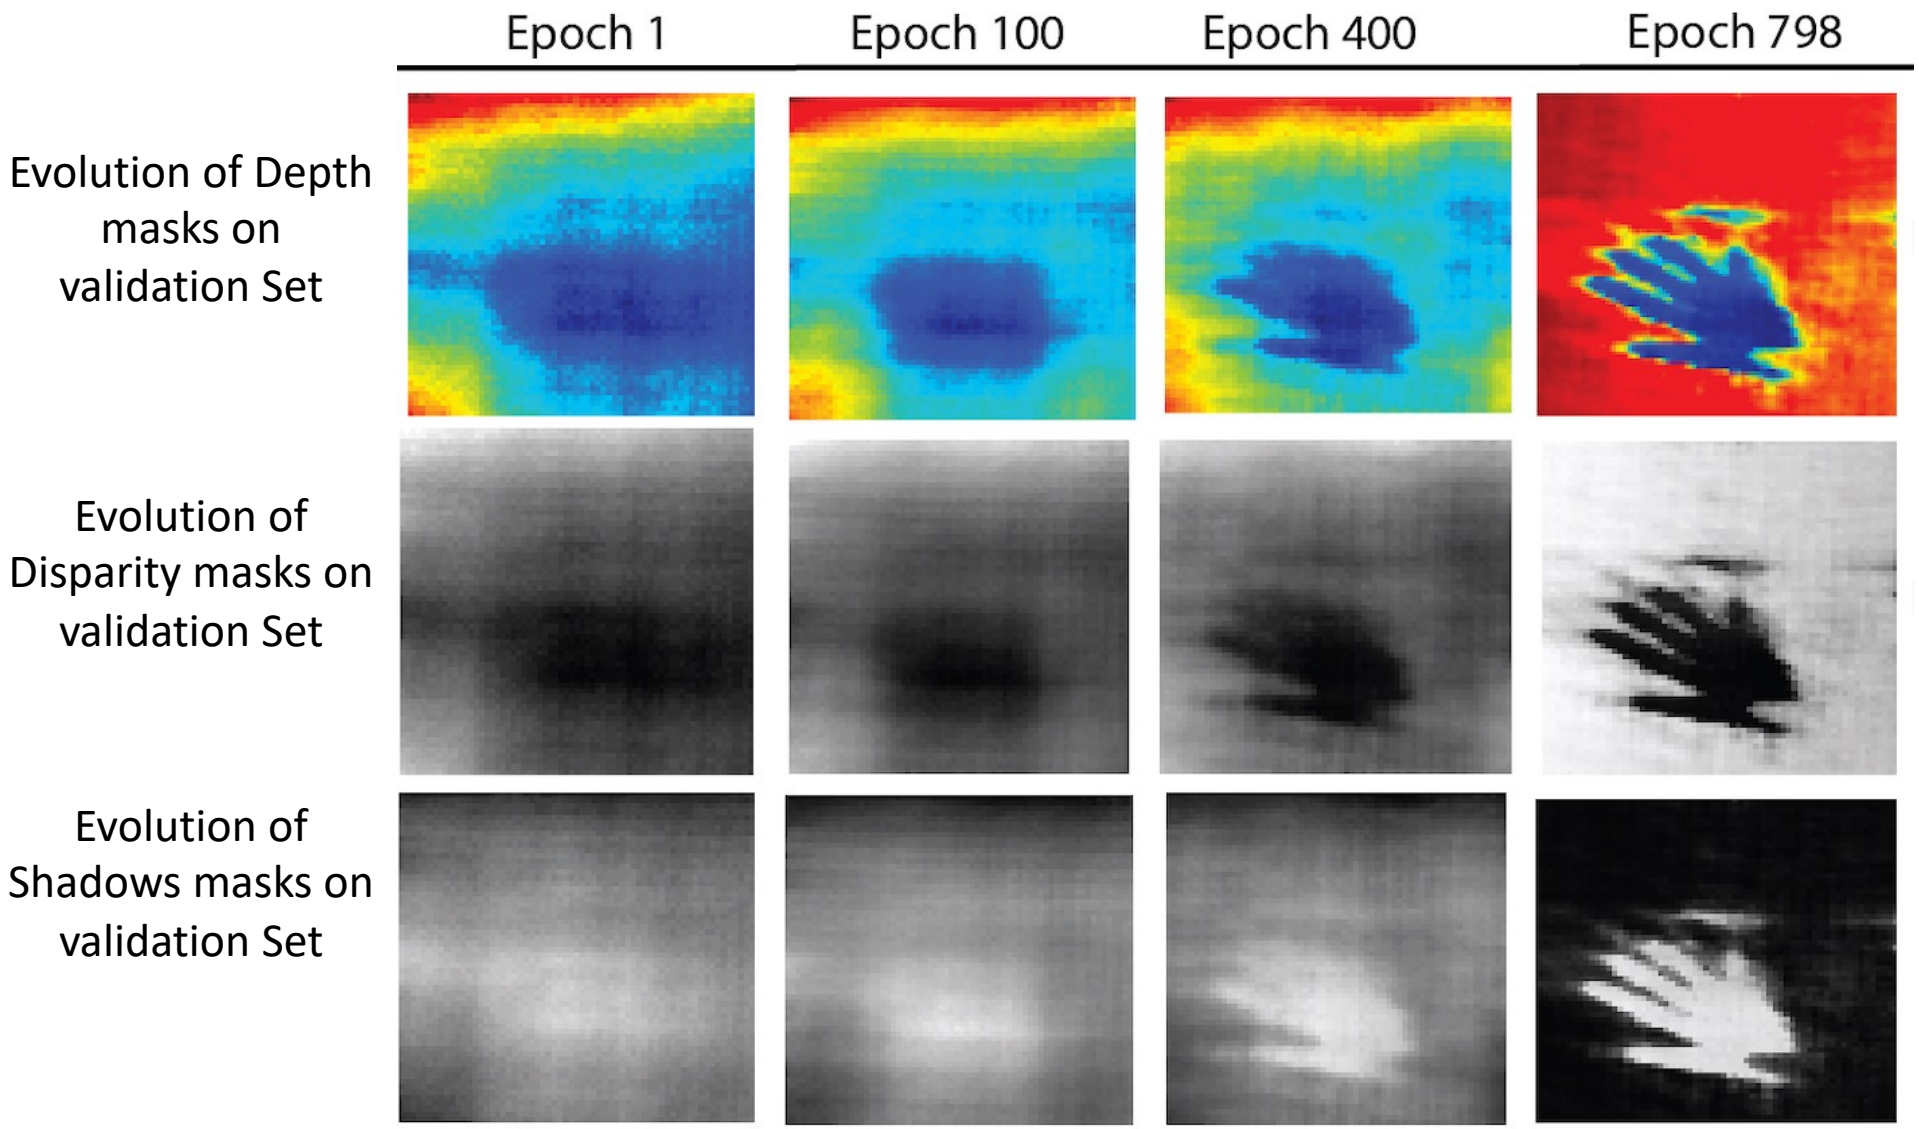

Validation Mask

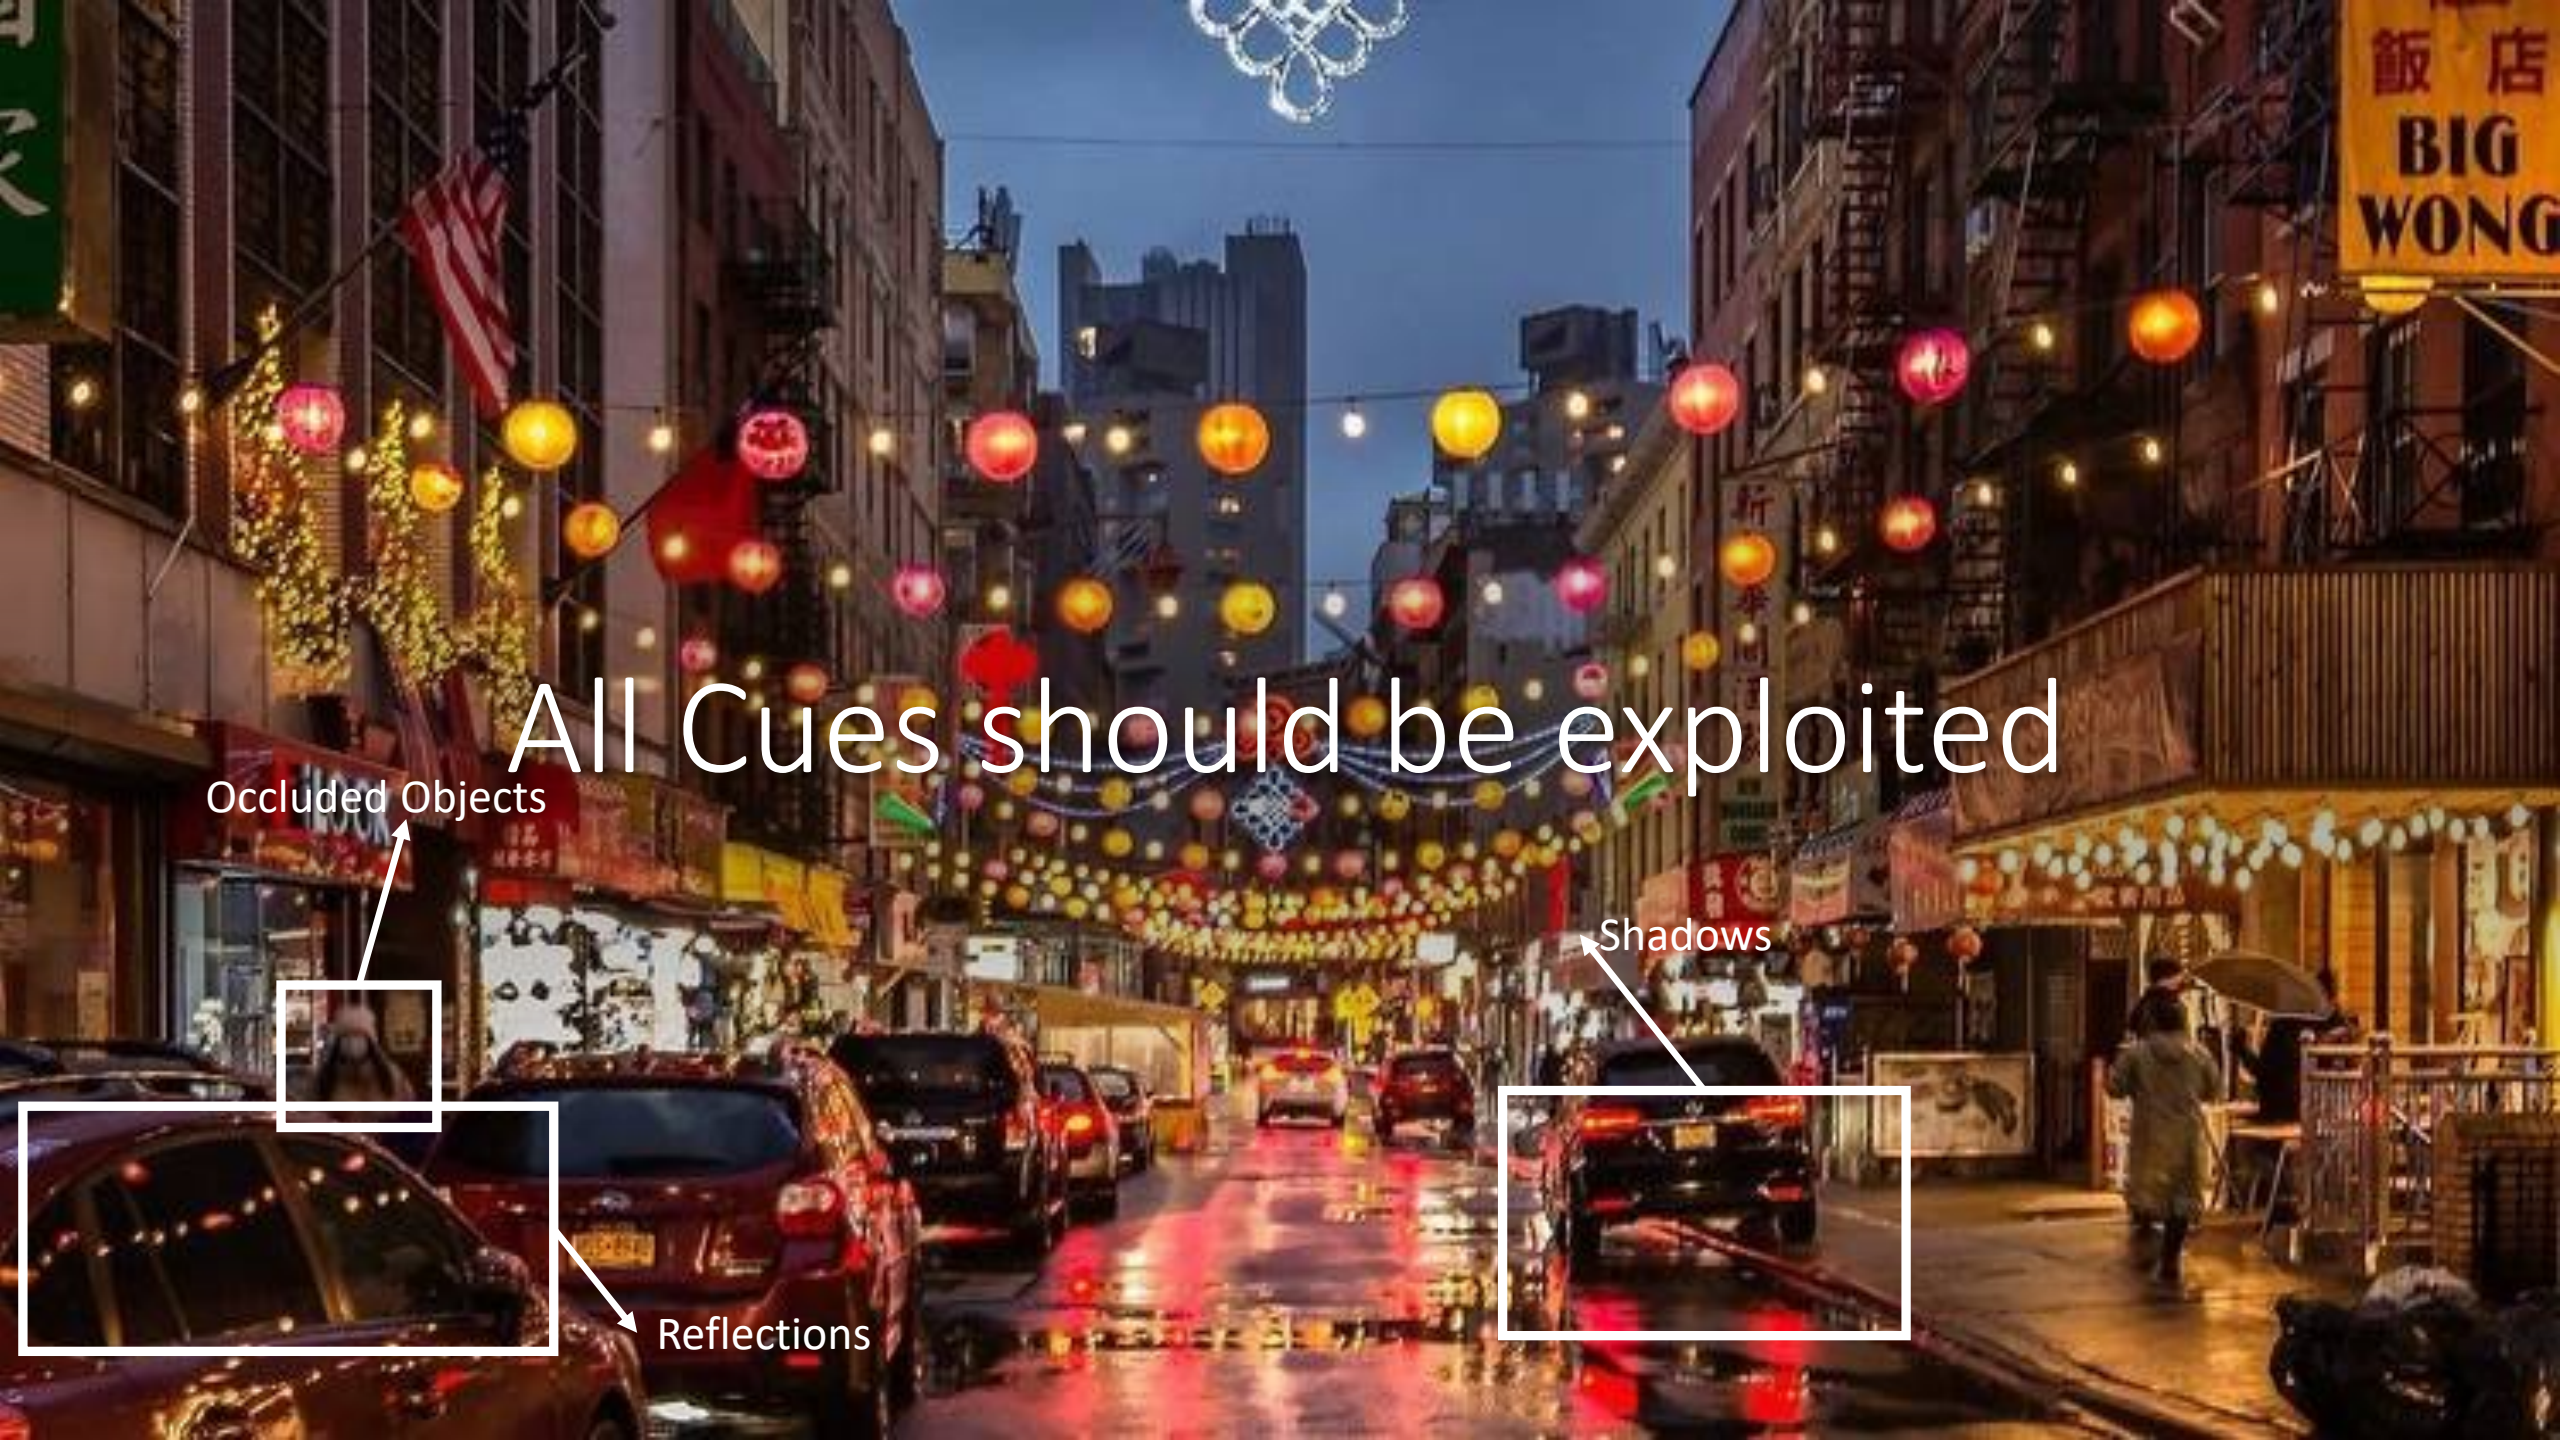
All Cues should be exploited
Occluded Objects
Shadows
Reflections
BIG
WONG

1000ft view: Exploiting All Cues

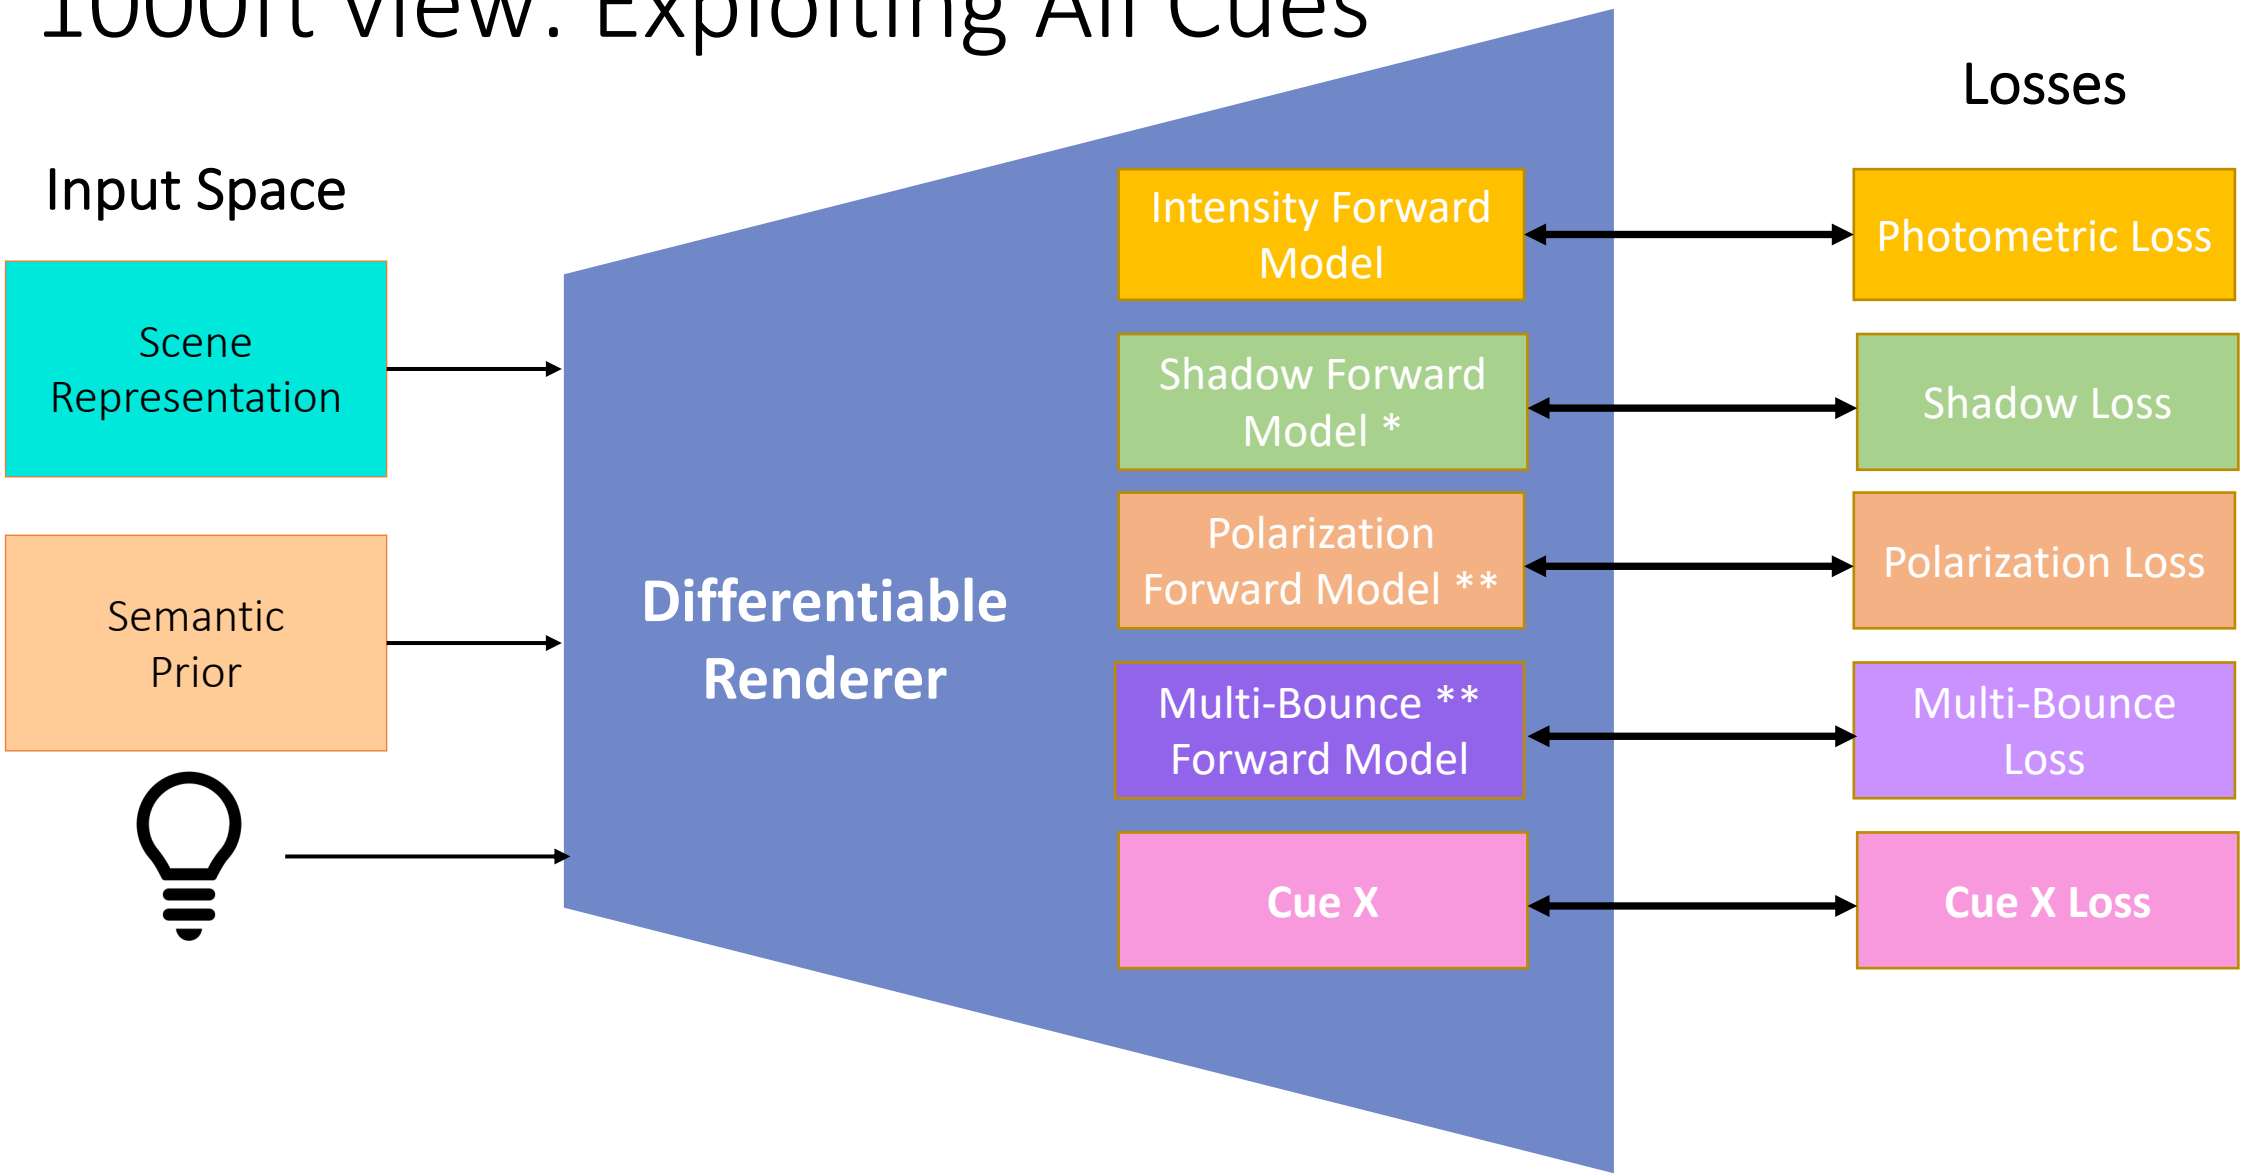

Tiwary, Klinghoffer*, and Raskar; Towards Learning Neural Representations from Shadows, *ECCV 2022*
**ORCA: Objects as Radiance Field Cameras, *Upcoming*

1000ft view: Exploiting All Cues

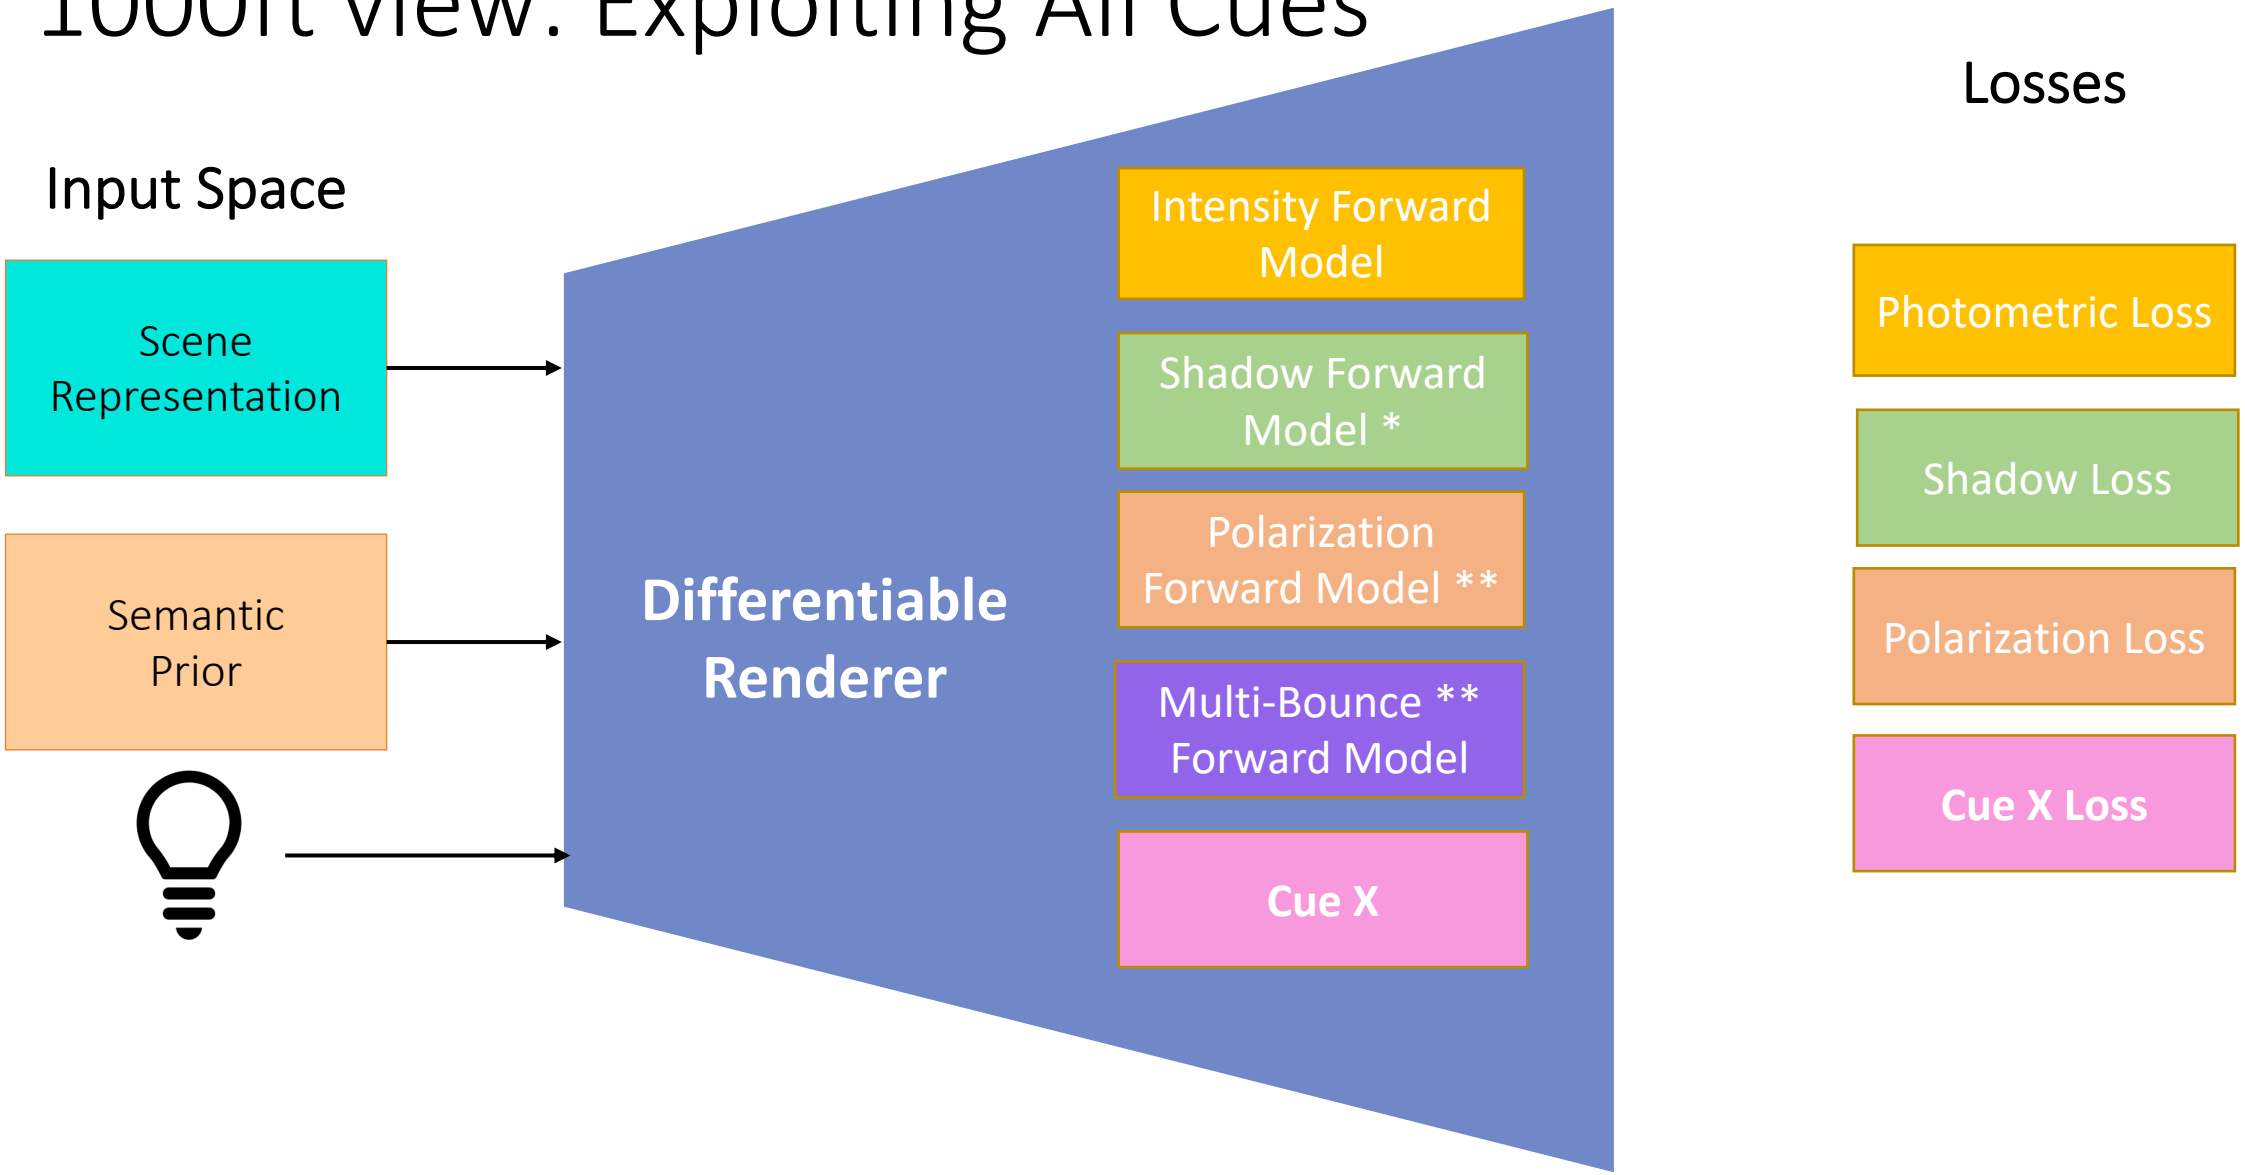

Tiwary, Klinghoffer*, and Raskar; Towards Learning Neural Representations from Shadows, *ECCV 2022*
**ORCA: Objects as Radiance Field Cameras, *Upcoming*

Summary

Ignore Shadows

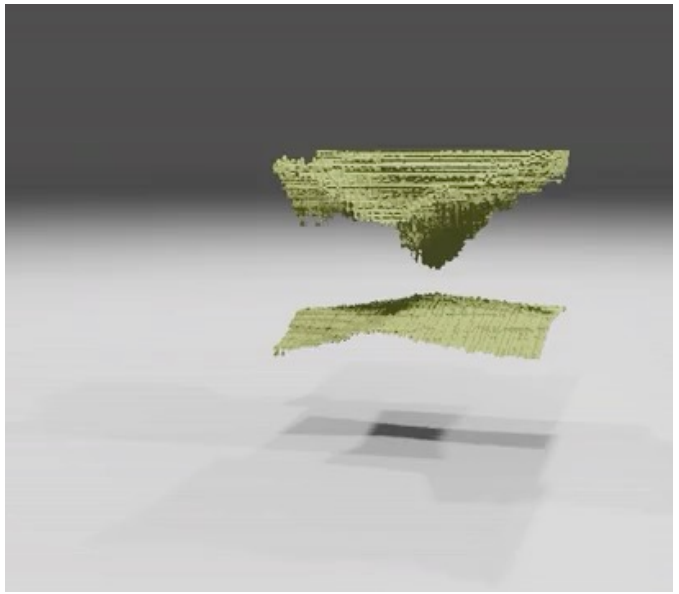

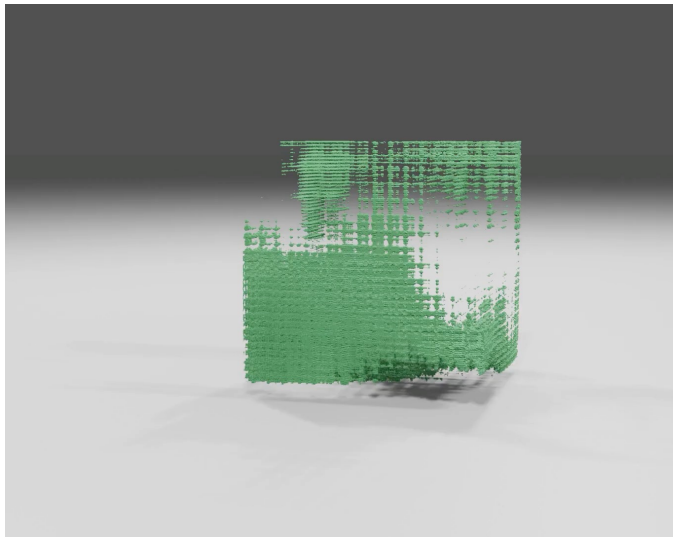

Exploit Shadows

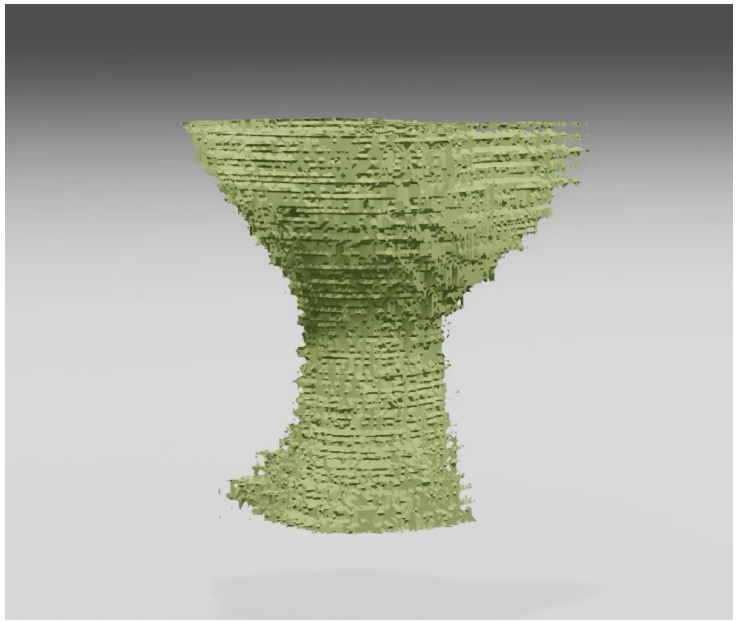

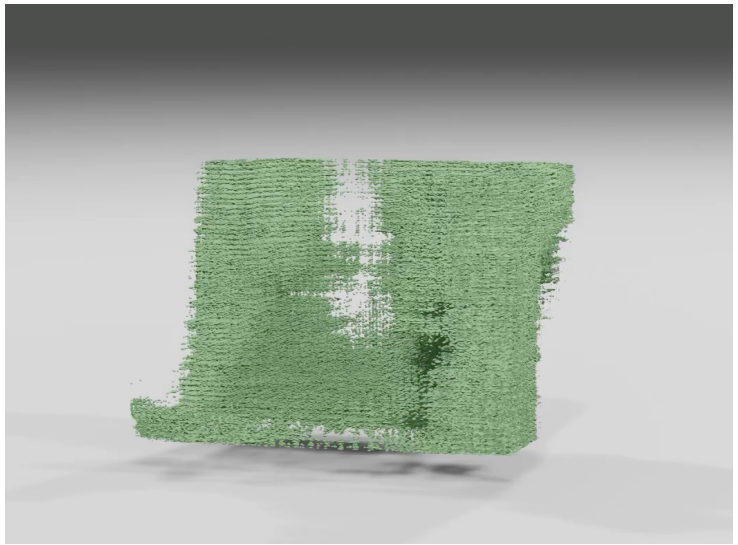

Can recover Hidden Geometry

Differentiable Shadow Forward Model

Recover Texture less objects