A Unified Spatial-Angular Structured Light for Single-View Acquisition of Shape and Reflectance

Xianmin Xu[1] Yuxin Lin[1] Haoyang Zhou[1] Chong Zeng[1] Yaxin Yu[1] Kun Zhou[1,2] Hongzhi Wu[1]

[1]State Key Lab of CAD&CG, Zhejiang University

[2]ZJU-FaceUnity Joint Lab of Intelligent Graphics

CVPR2023 TAG: TUE-AM-020

Overview

1. Introduction
2. Related Work
3. Hardware Prototype
4. Depth Acquisition
5. Reflectance Capture
6. Results
7. Limitations and Future Work

1 Introduction

Motivation

- Active lighting in high-quality acquisition

Active Lighting

Shape

HandySCAN 3D

Material

Lightstage [Kang et al. 2019]

Shape & Material

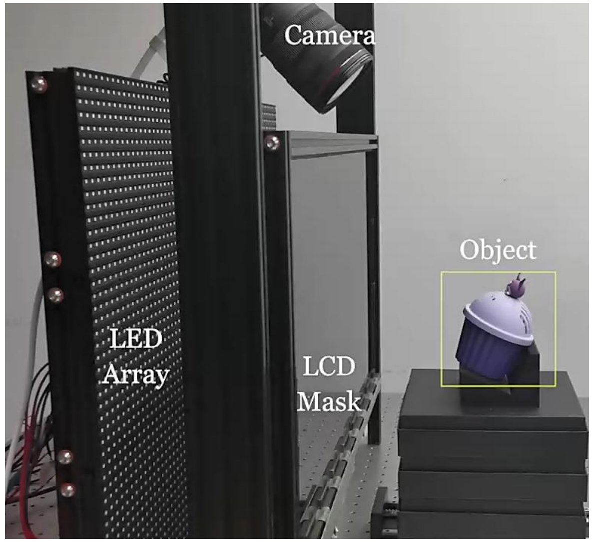

Our Lightweight Setup

1 Introduction

Acquisition Setup

- A novel lightweight acquisition setup acts as a restricted lightfield projector

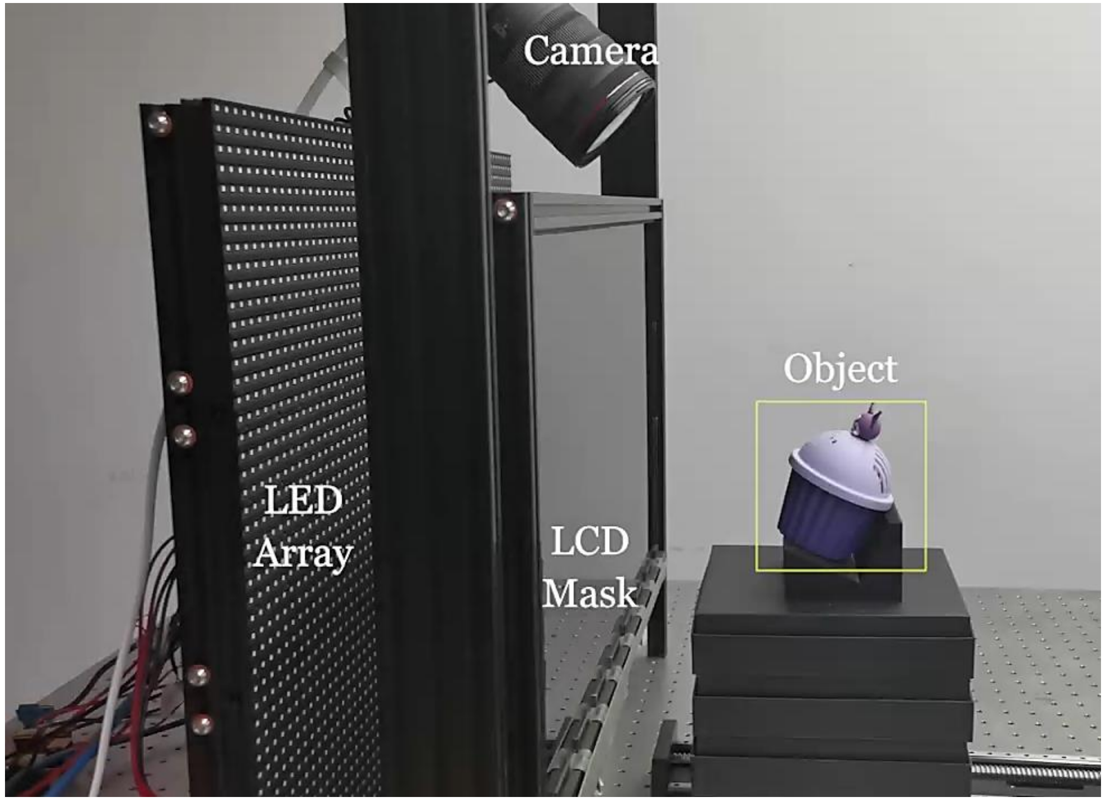

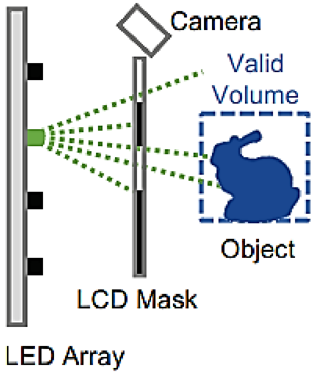

Depth Acquisition

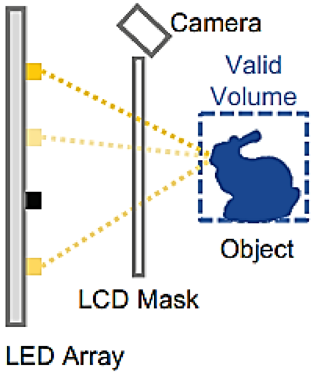

Reflectance Acquisition

1 Introduction

Pipeline

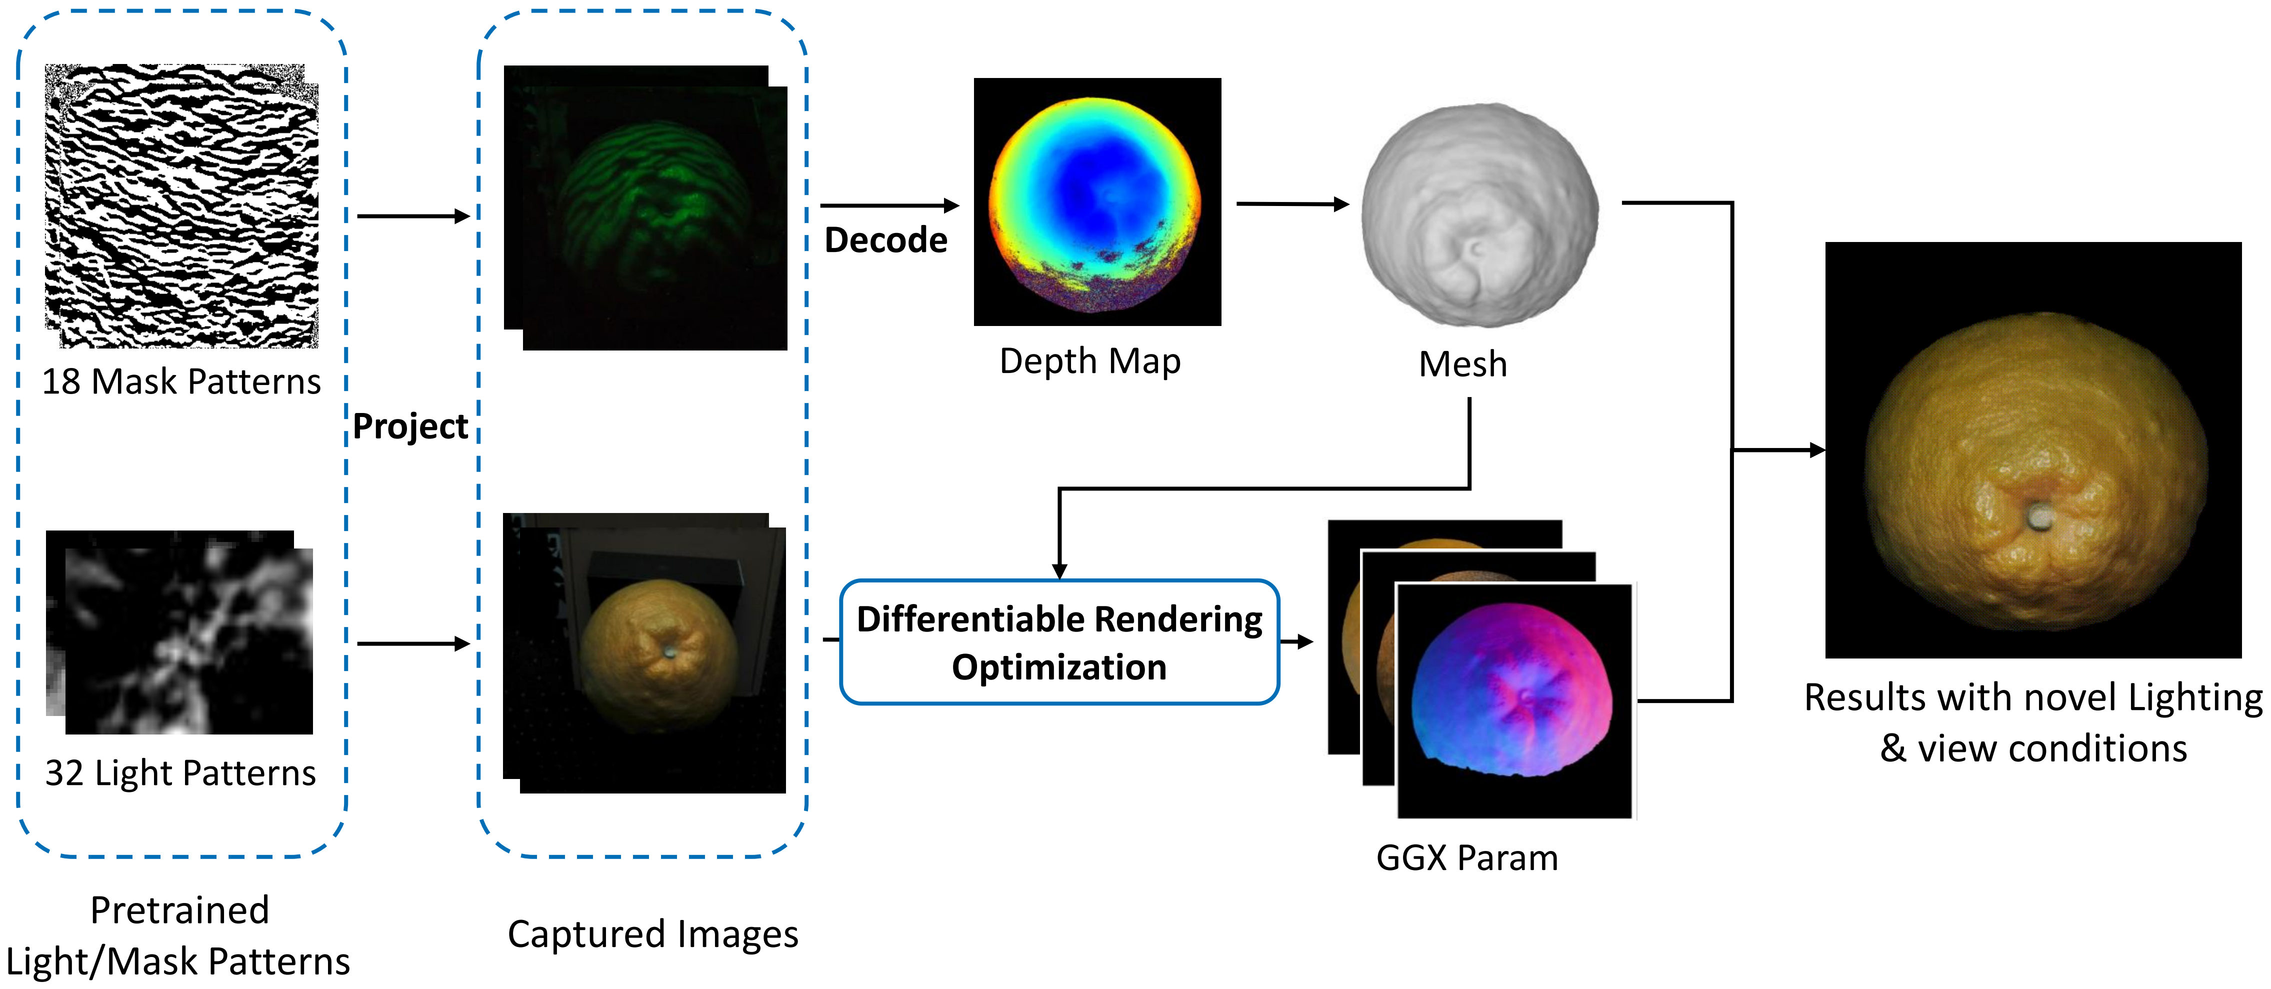

2 Related Work

Multi-Pattern Structured-Light

- Predetermined Patterns vs Optimized Patterns

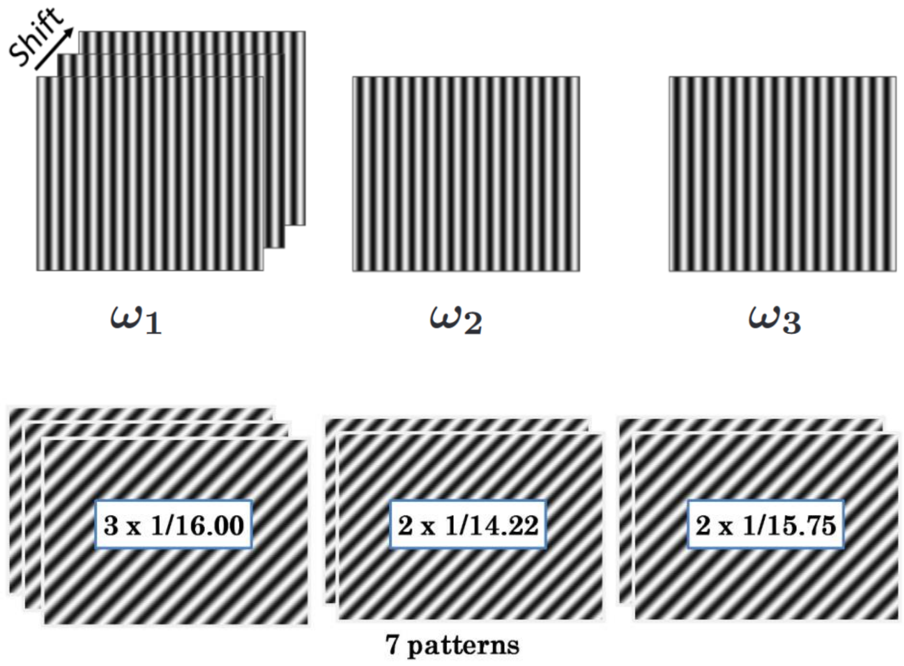

MPS [Gupta et al. 2012] & EPS [Moreno et al. 2015]

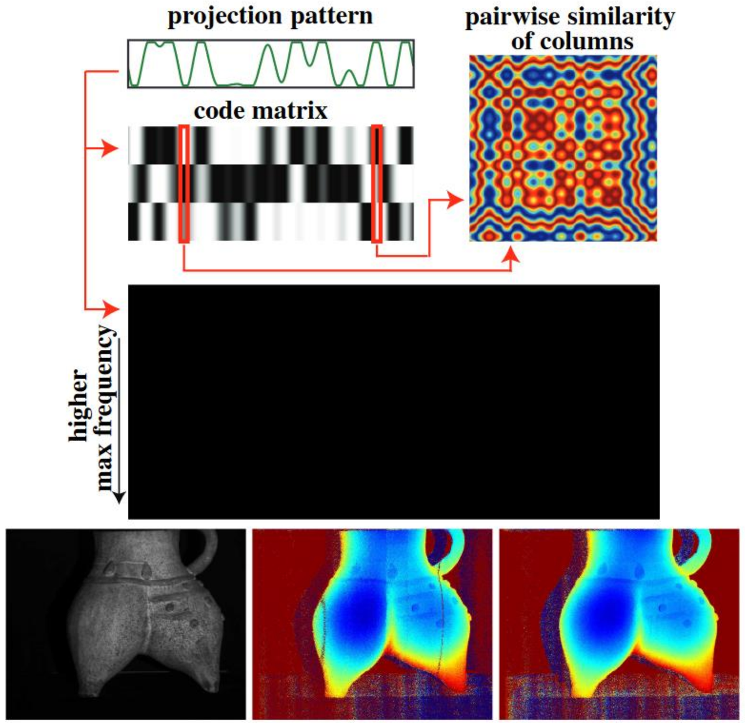

A la Carte [Mirdehghan et al. 2018]

2 Related Work

Illumination Multiplexing

- Differentiable Acquisition [Kang et al. 2018; Kang et al. 2019; Ma et al. 2021]
  - Obtain high-fidelity reflectance results
  - Limitations: visibility not considered

Light Cube [Kang et al. 2019]

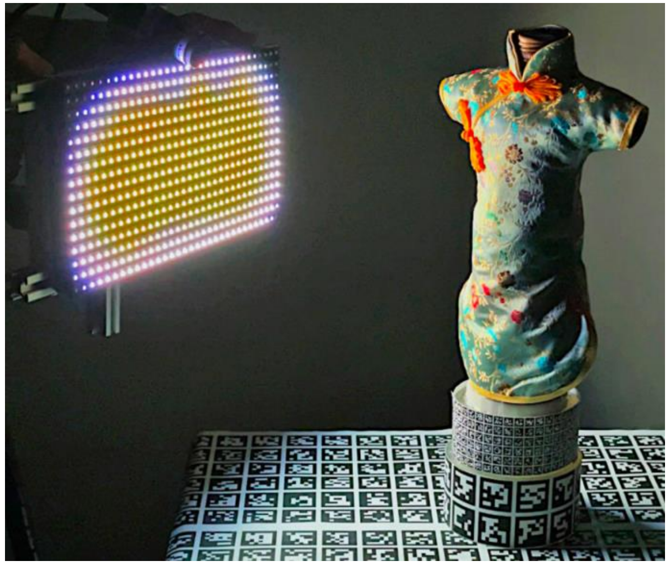

MatScan [Ma et al. 2021]

3 Hardware Prototype

Our Acquisition Setup

- RGB LEDs: 64×48 = 3, 072
- LCD Mask: resolution 1920×1080, size 59.8cm×33.6cm
- 45MP Canon EOS R5 Camera
- Valid Volume: 15cm×15cm ×15cm

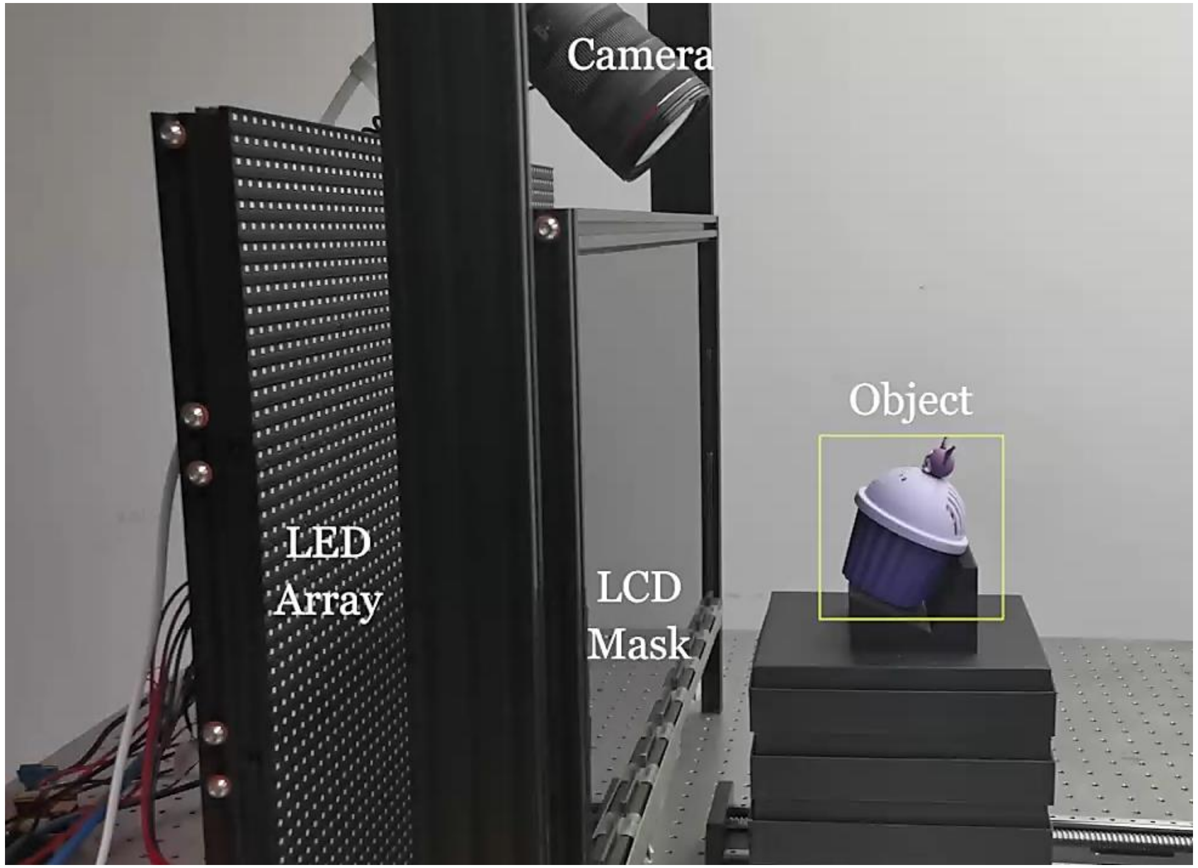

Our Lightweight Acquisition Setup

3 Hardware Prototype

Differentiable Calibration

- Calibrate each LED **individually** (position, normal, size)
- Training data: correspondences between mask & image
- Optimize LED position and transformation of mask plane by **stochastic gradient descent** (SGD)

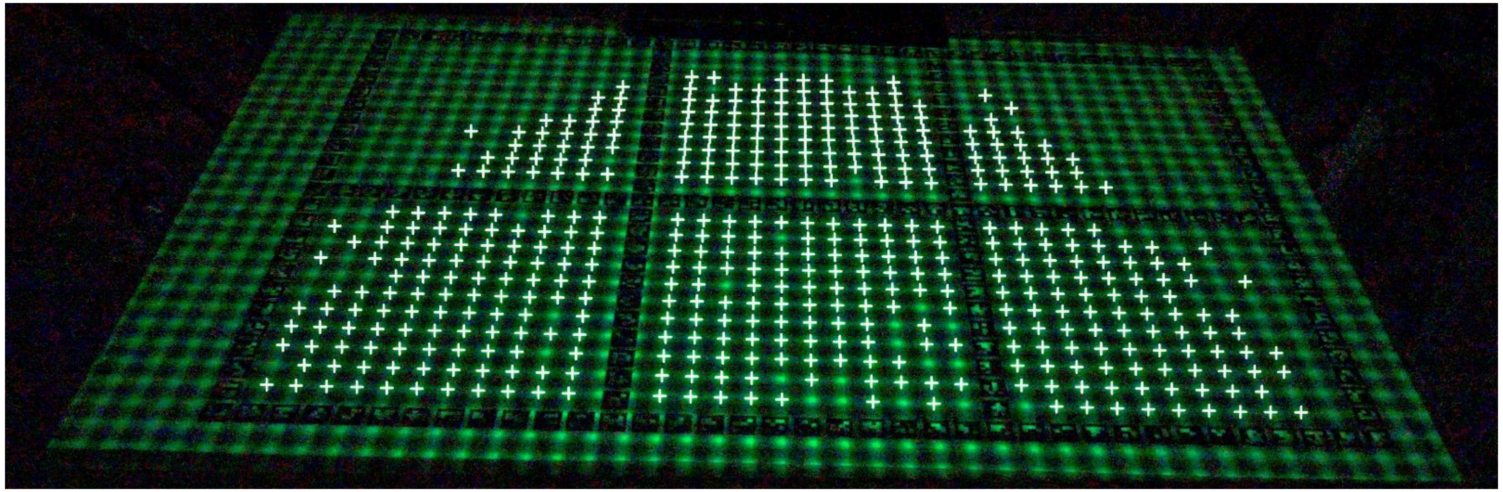

A Dense Set of Crosses on Board

Sample Candidates along Camera Ray

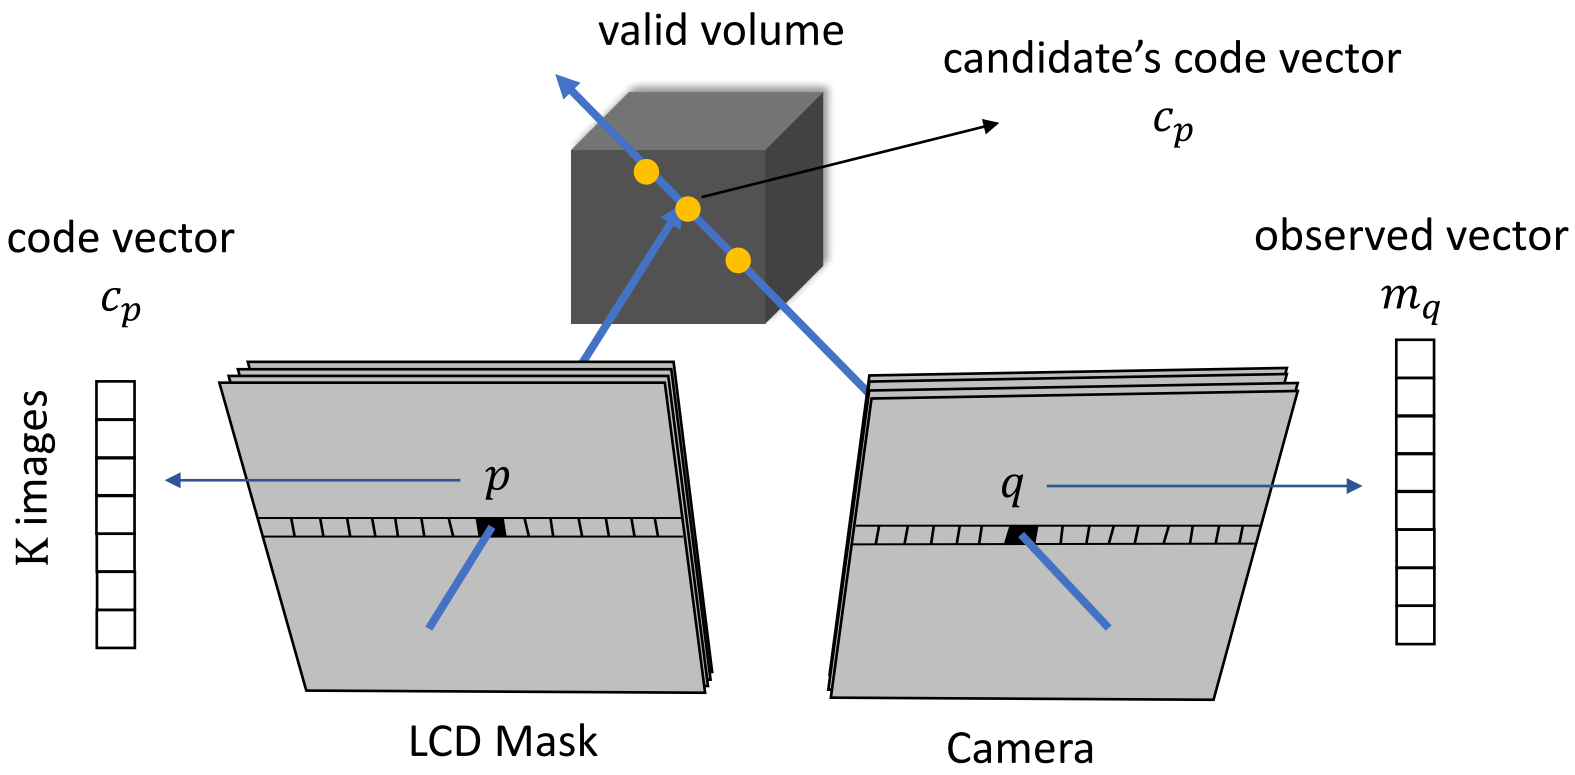

The number of mask pixels is limited, but we can sample any number of candidates.

4 Depth Acquisition

Area Light Convolution

- The area light will perform **convolution operation** on the **binary** mask pattern.
- The code will change **continuously** along the camera ray.

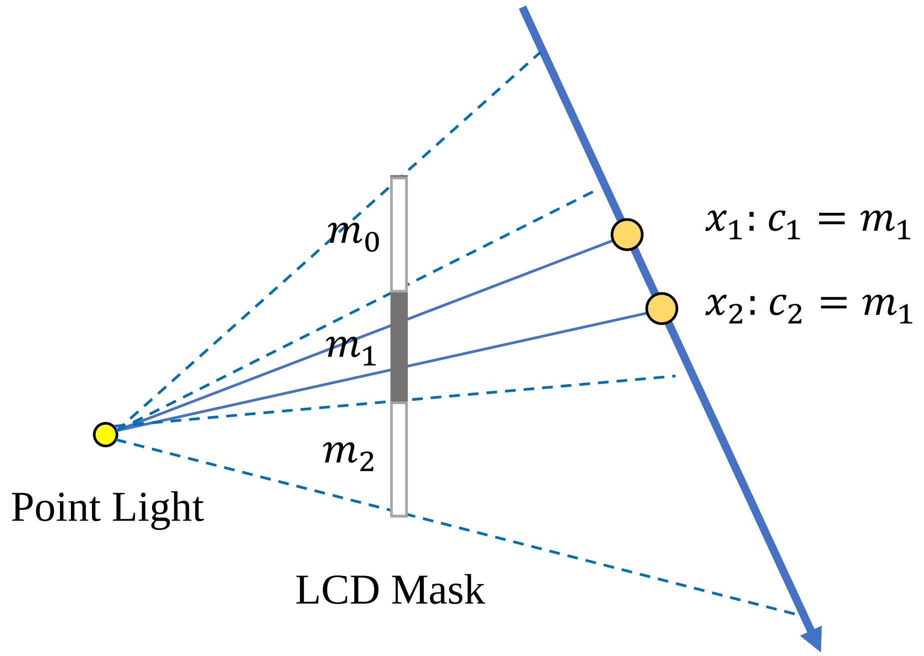

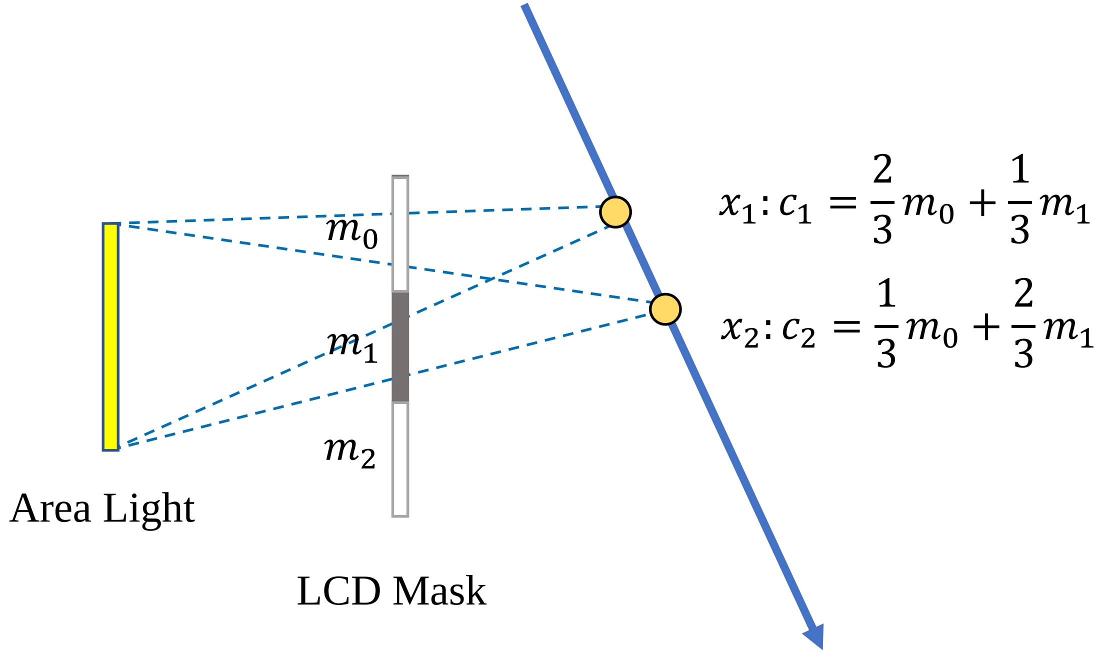

4 Depth Acquisition

Depth Decoding

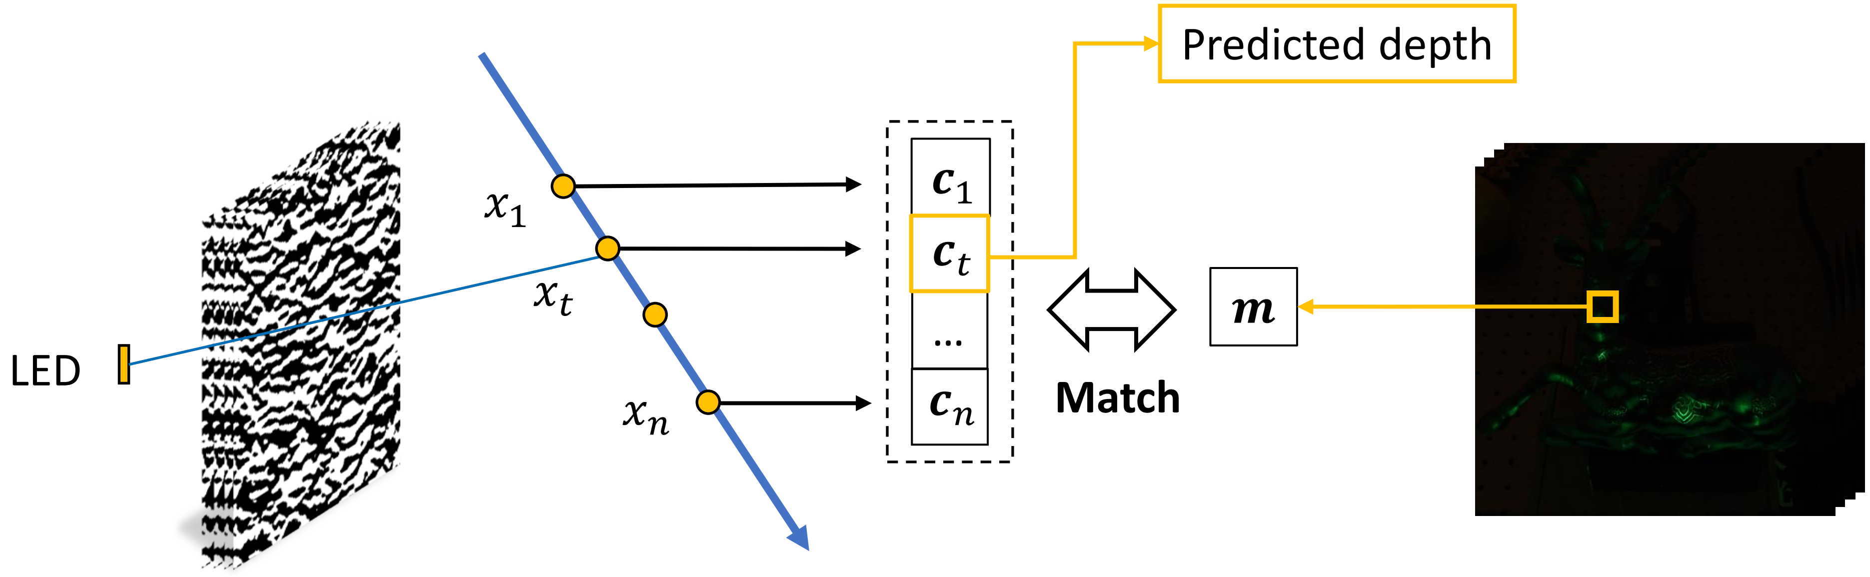

M Mask Patterns

M Images

Match: Find highest $\mathrm{ZNCC}(\boldsymbol{c}, \boldsymbol{m}) = \frac{\boldsymbol{c} - \mathrm{mean}(\boldsymbol{c})}{||\boldsymbol{c} - \mathrm{mean}(\boldsymbol{c})||} \cdot \frac{\boldsymbol{m} - \mathrm{mean}(\boldsymbol{m})}{||\boldsymbol{m} - \mathrm{mean}(\boldsymbol{m})||}$

Mask Pattern Pretraining

- Cast to a multi-class classification problem

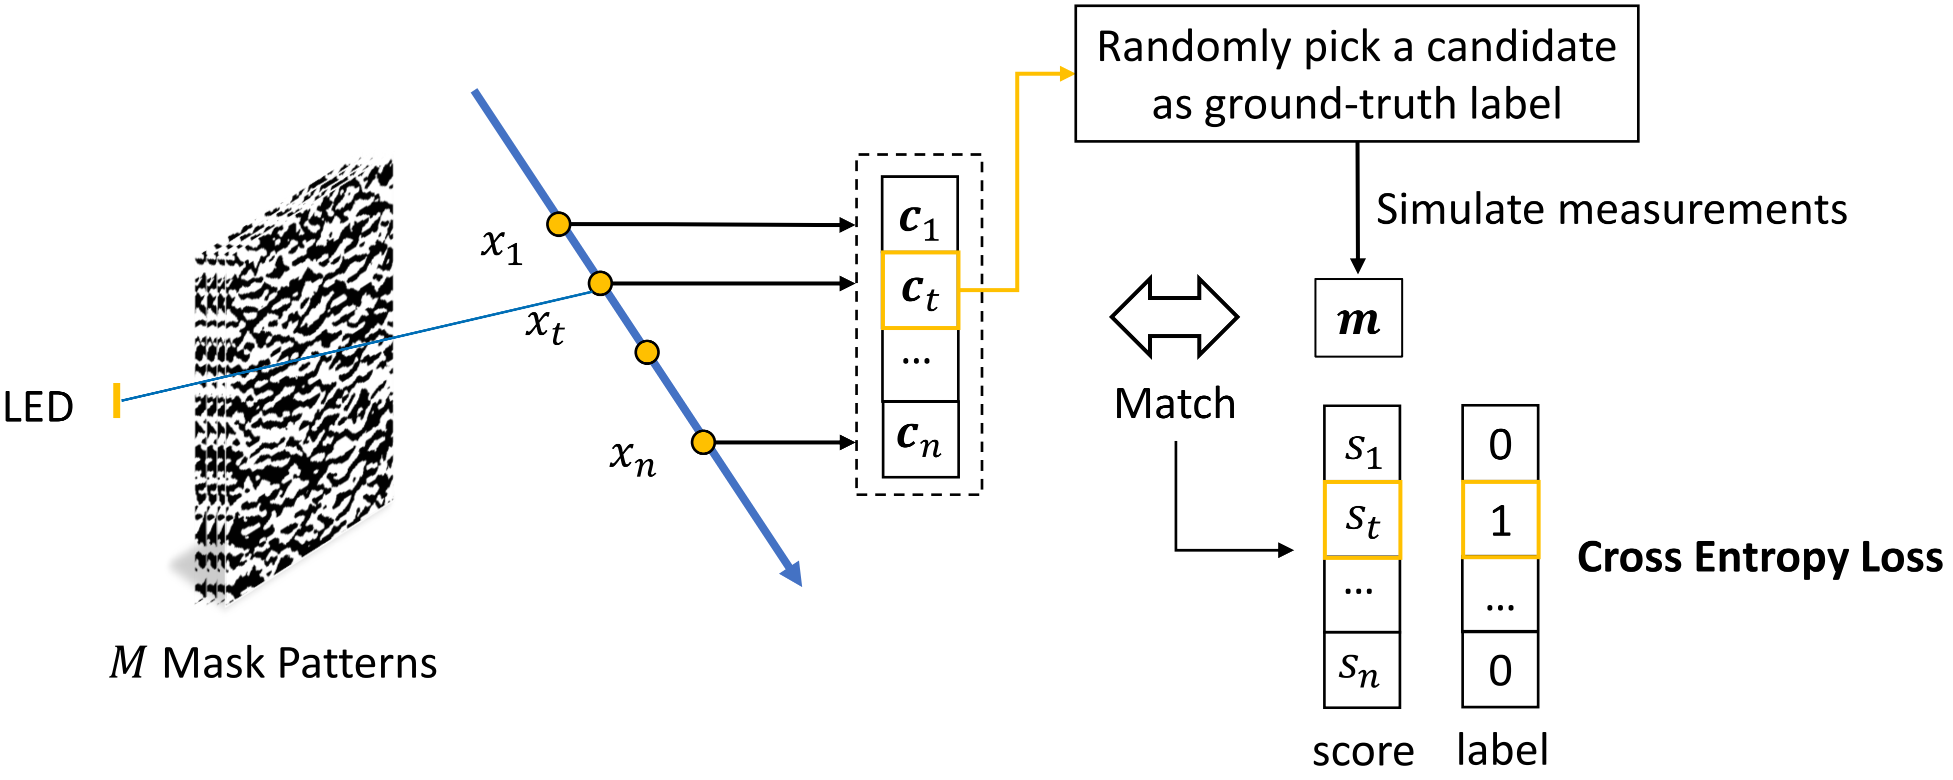

4 Depth Acquisition

Multiple LEDs

- Simply concatenate the code and the measurements
- Geometric accuracy and completeness can be further improved

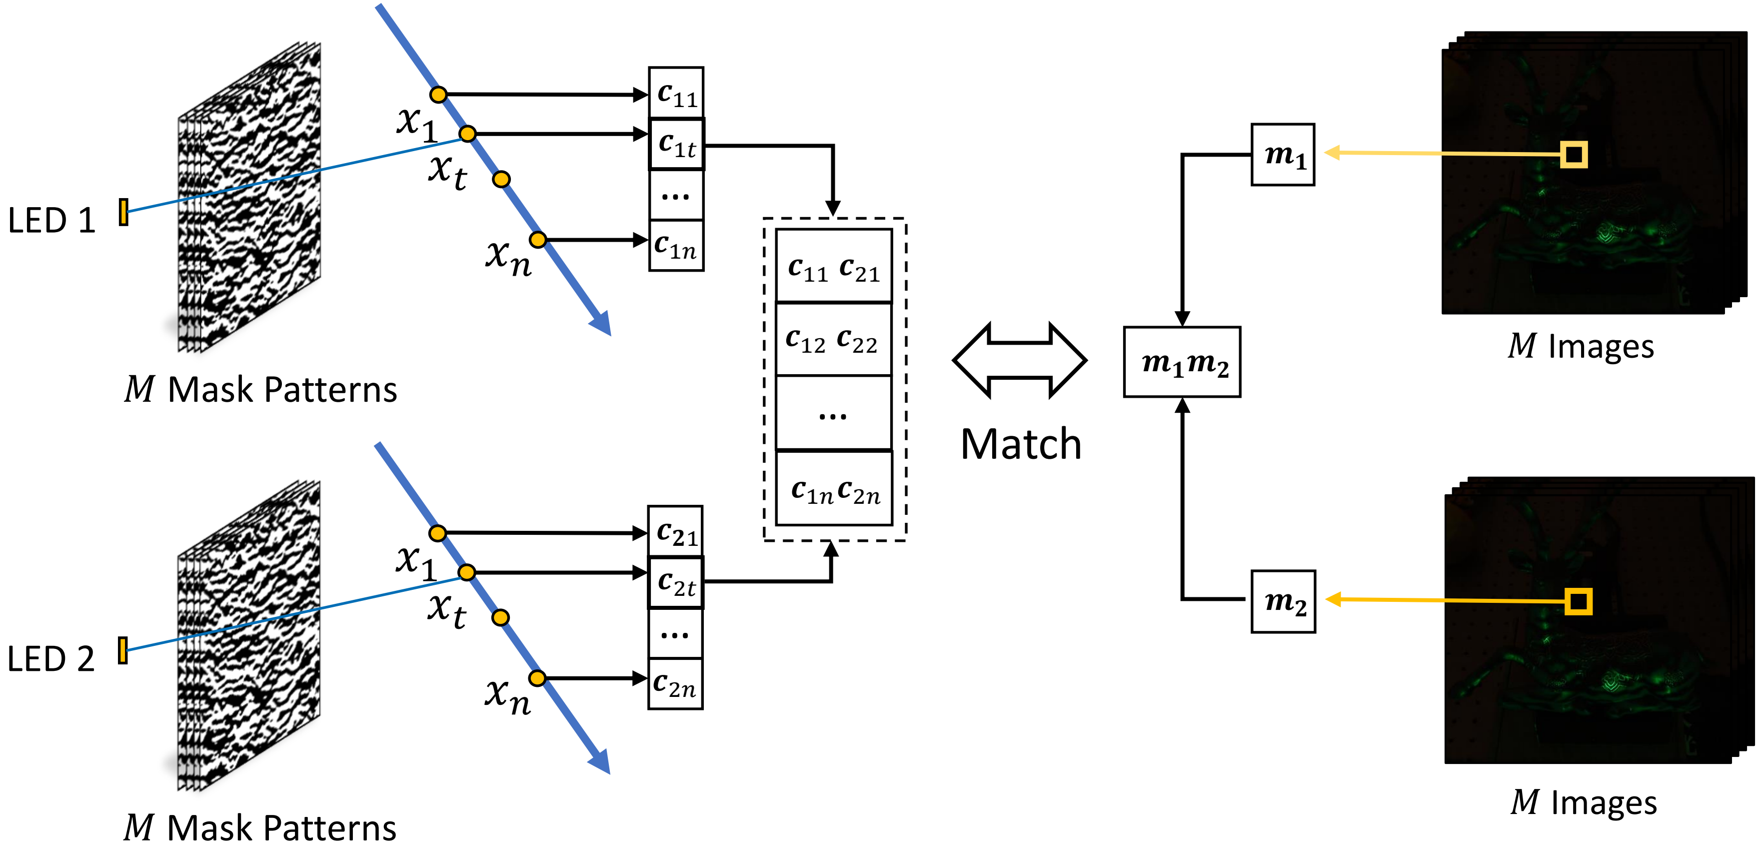

5 Reflectance Capture

Reflectance Pipeline

Light Pattern Training

- Train autoencoder same as [Kang et al. 2018]

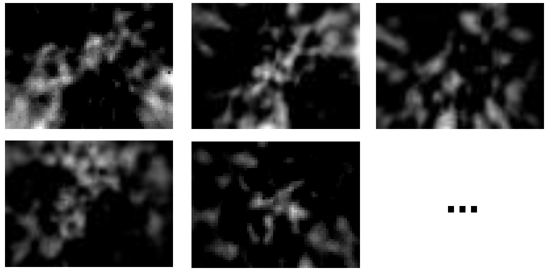

32 Light Patterns

Acquisition

- Set mask to transparent

32 Captured Images

Differentiable Rendering

- Discard the decoder
- Consider visibility

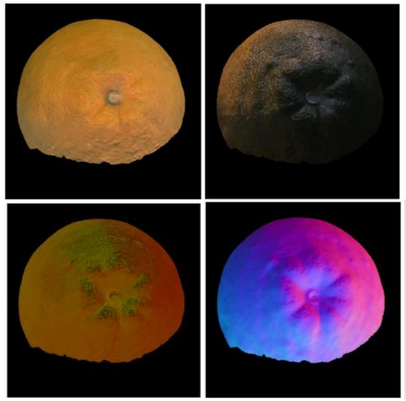

GGX Parameters

5 Reflectance Capture

Differentiable Optimization

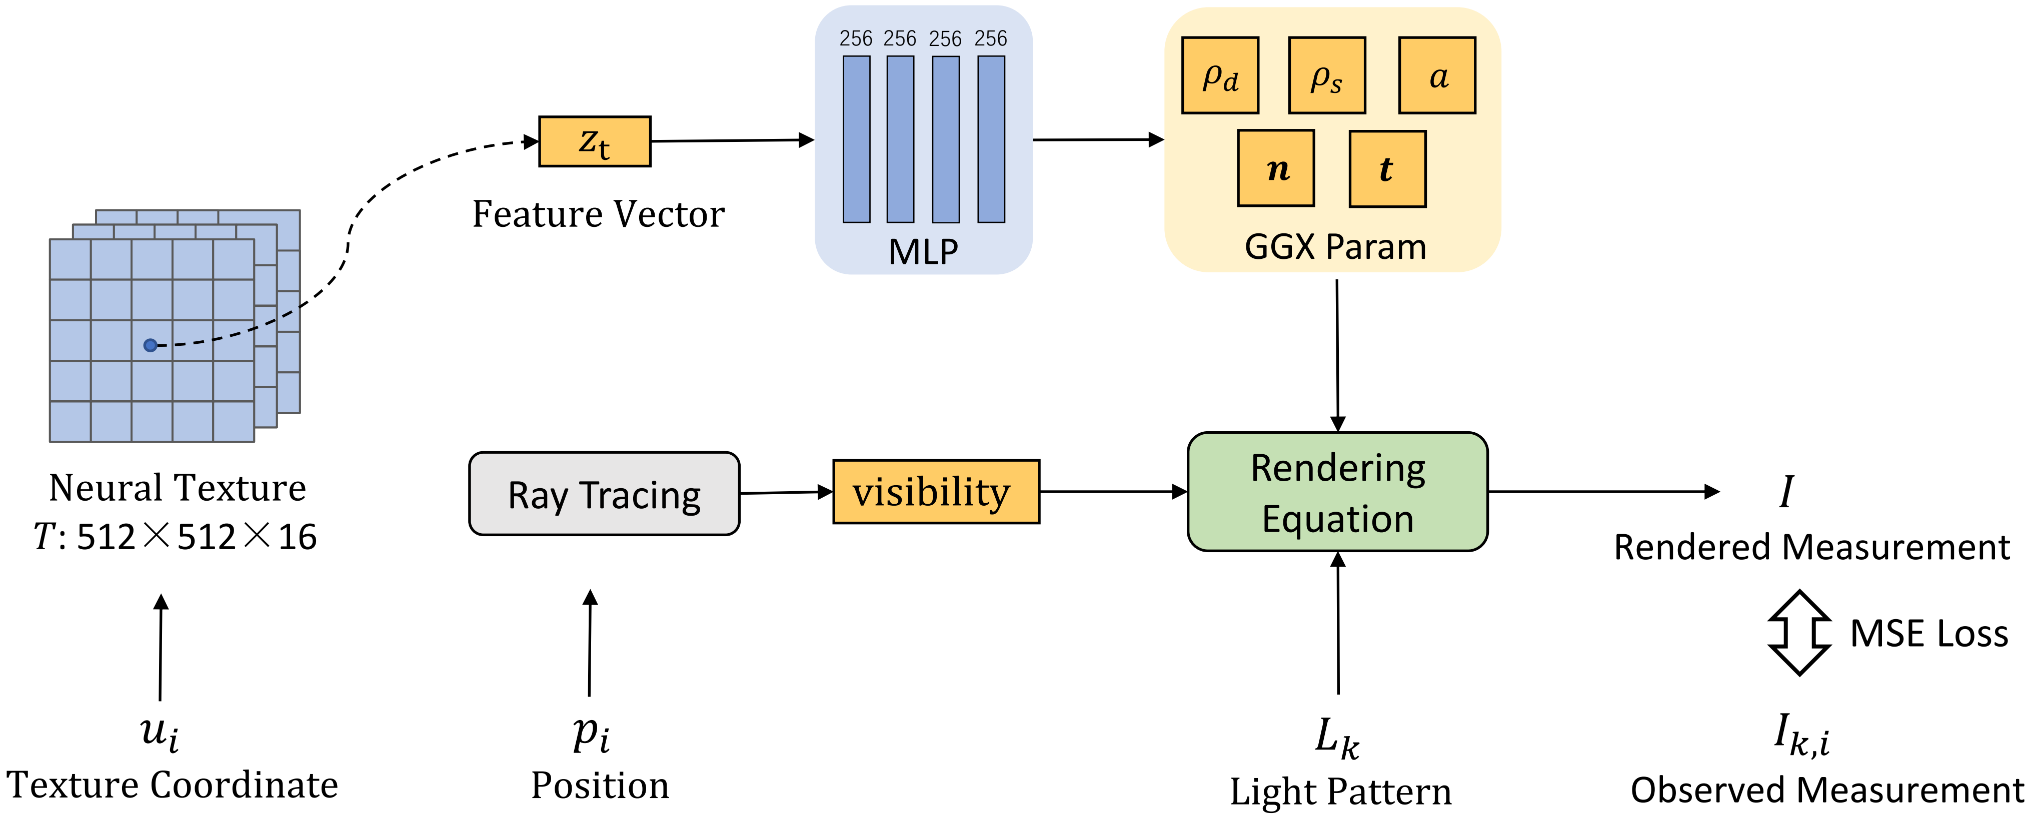

6 Results

Statistics

- Mask Pattern
  - Number: 18 for a single LED, 4 LEDs in total
  - Training: 15 minutes for a single LED
  - Acquisition: 24 minutes
  - Depth Decode 8 minutes
- Light Pattern
  - Number: 32
  - Acquisition: about 10 seconds
  - Reflectance Fine-tune: 1 hour

6 Results

Learned Patterns and Photos

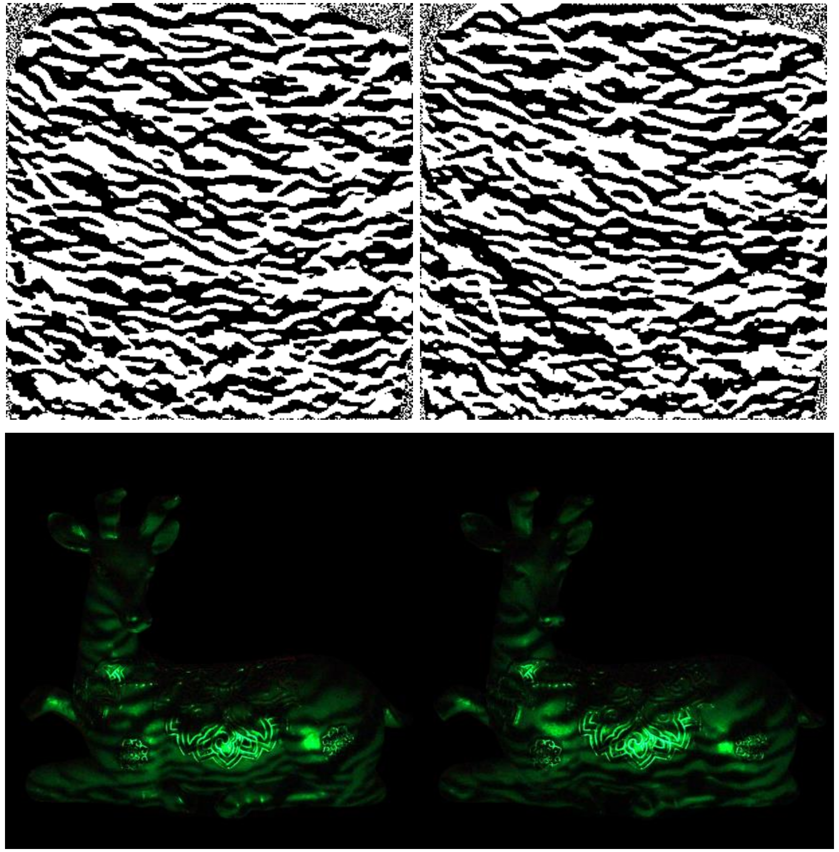

Mask Pattern

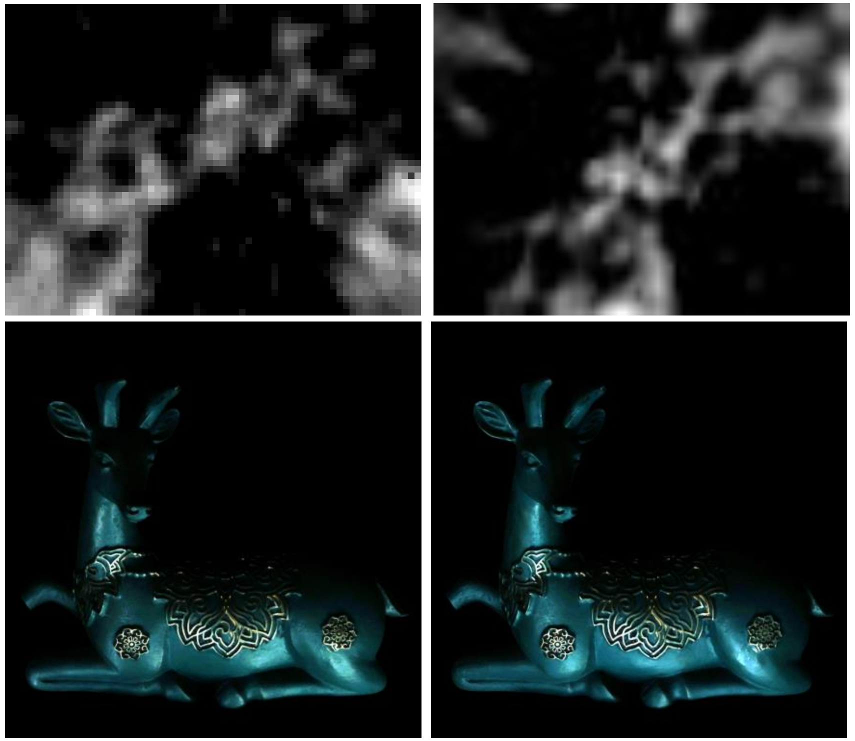

Light Pattern

6 Results

Depth Comparisons

- We achieve on a average geometric accuracy of **0.27mm**

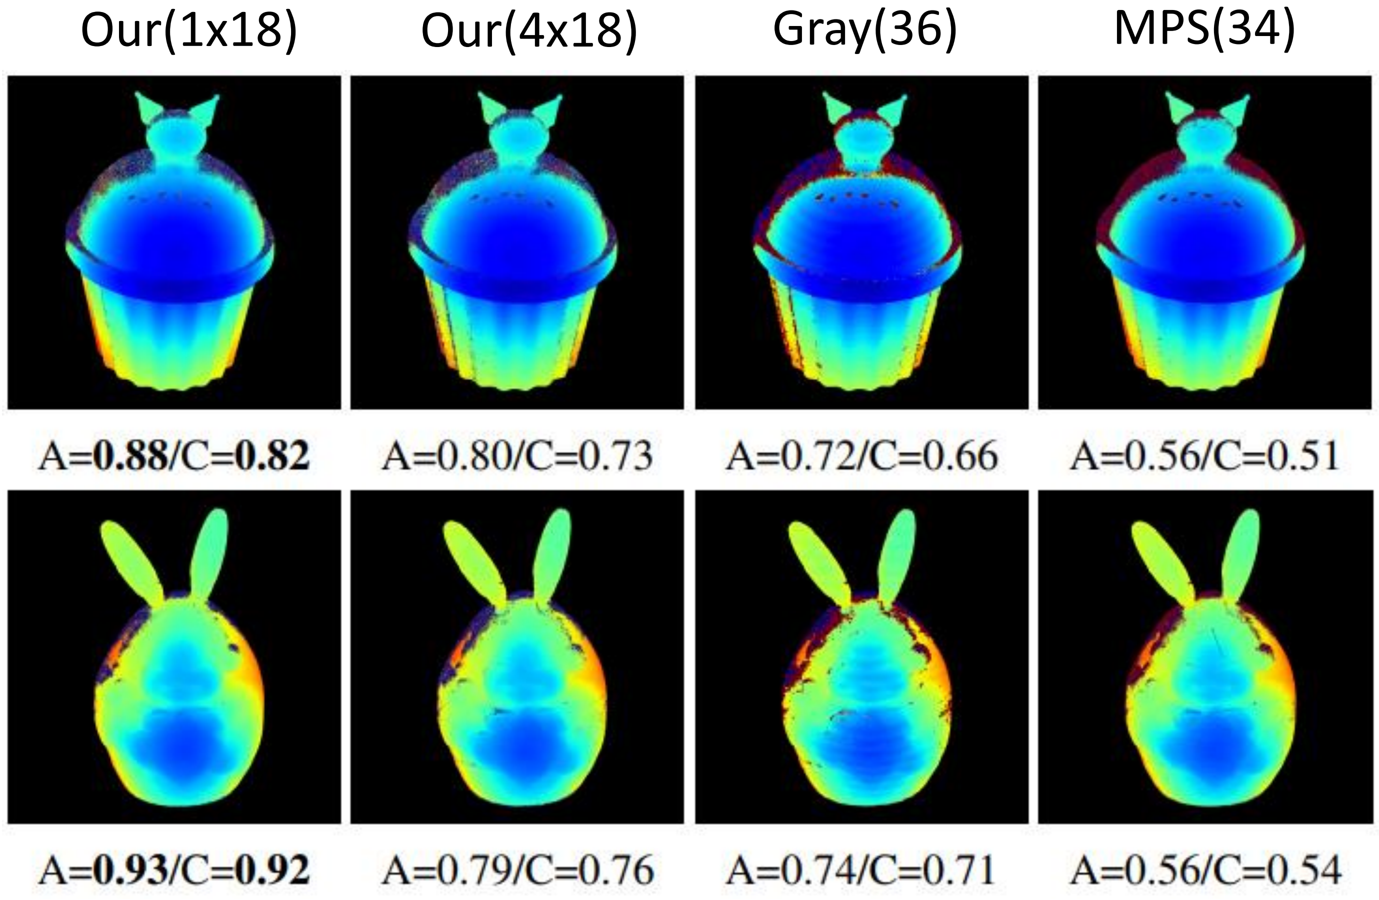

6 Results

Reflectance Comparisons

- Our results outperform state-of-the-art method [Kang et al. 2019]

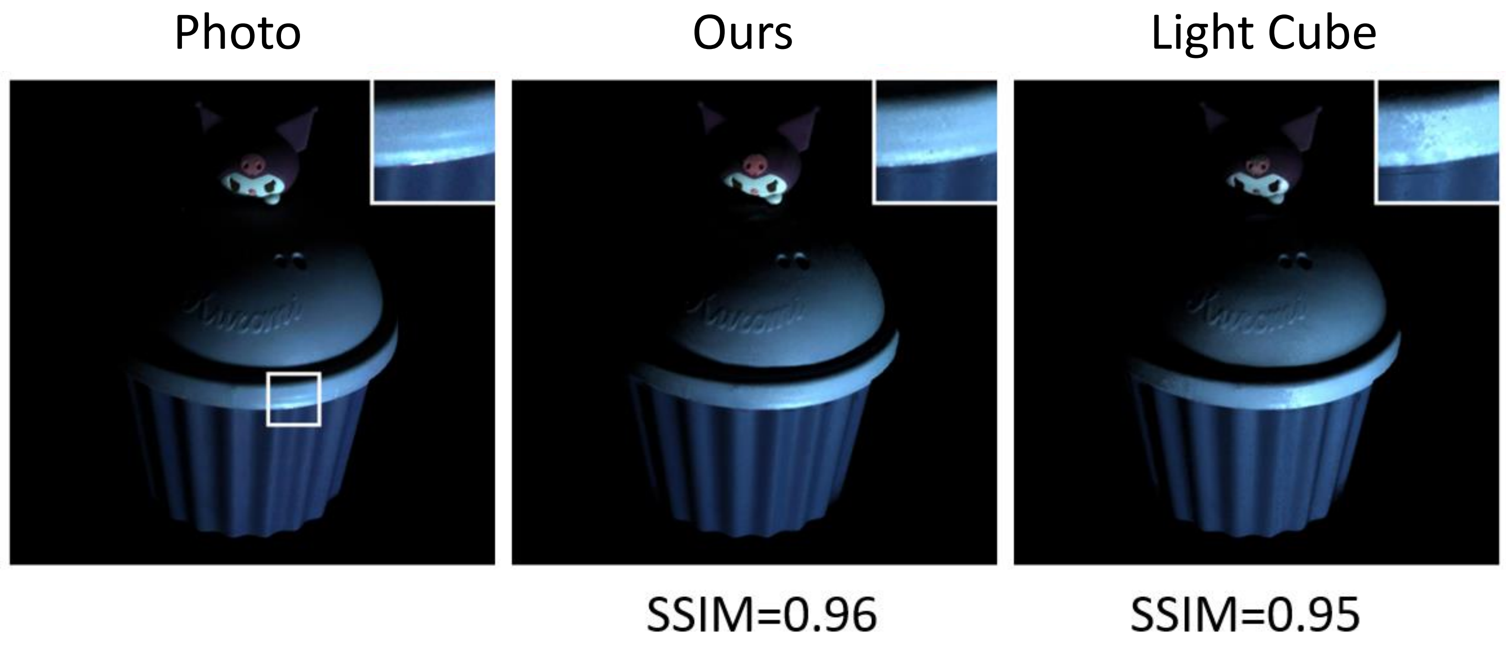

6 Results

Ablation Study

- Impact of **the number of mask patterns** over reconstructed depth

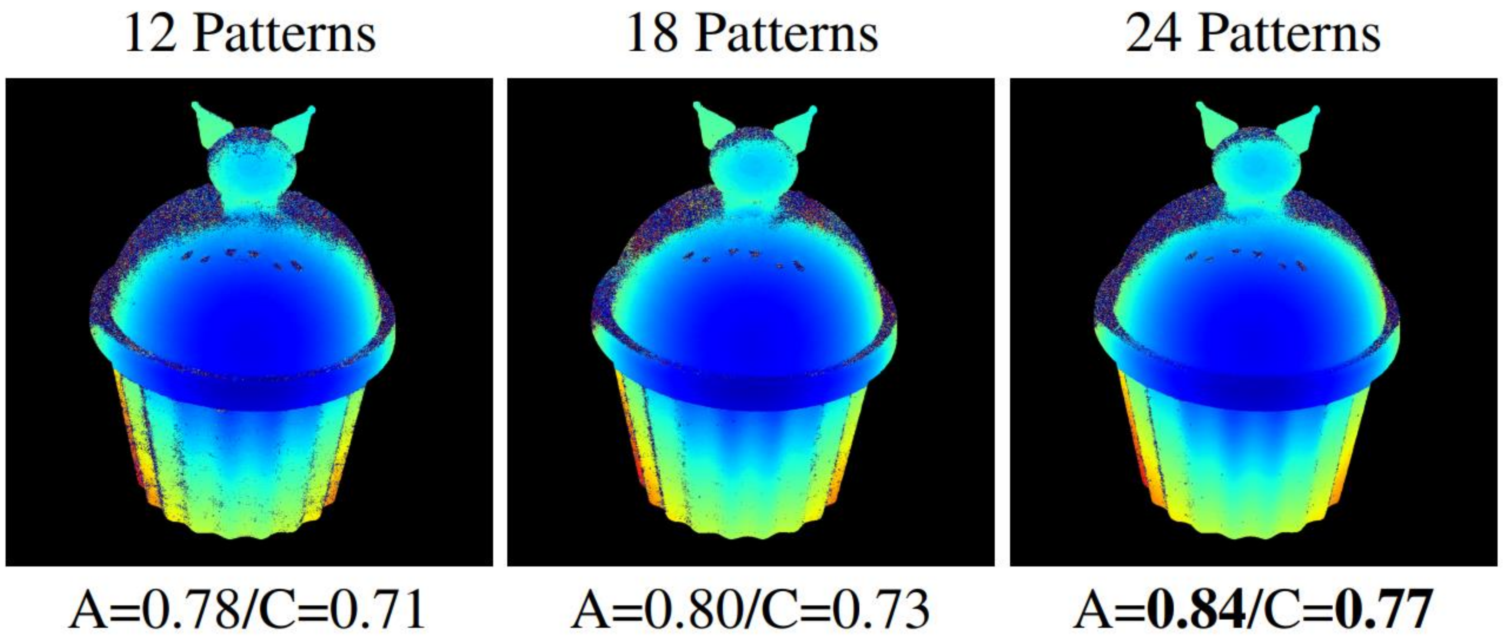

6 Results

Ablation Study

- Impact of **the number of LEDs** over reconstructed depths

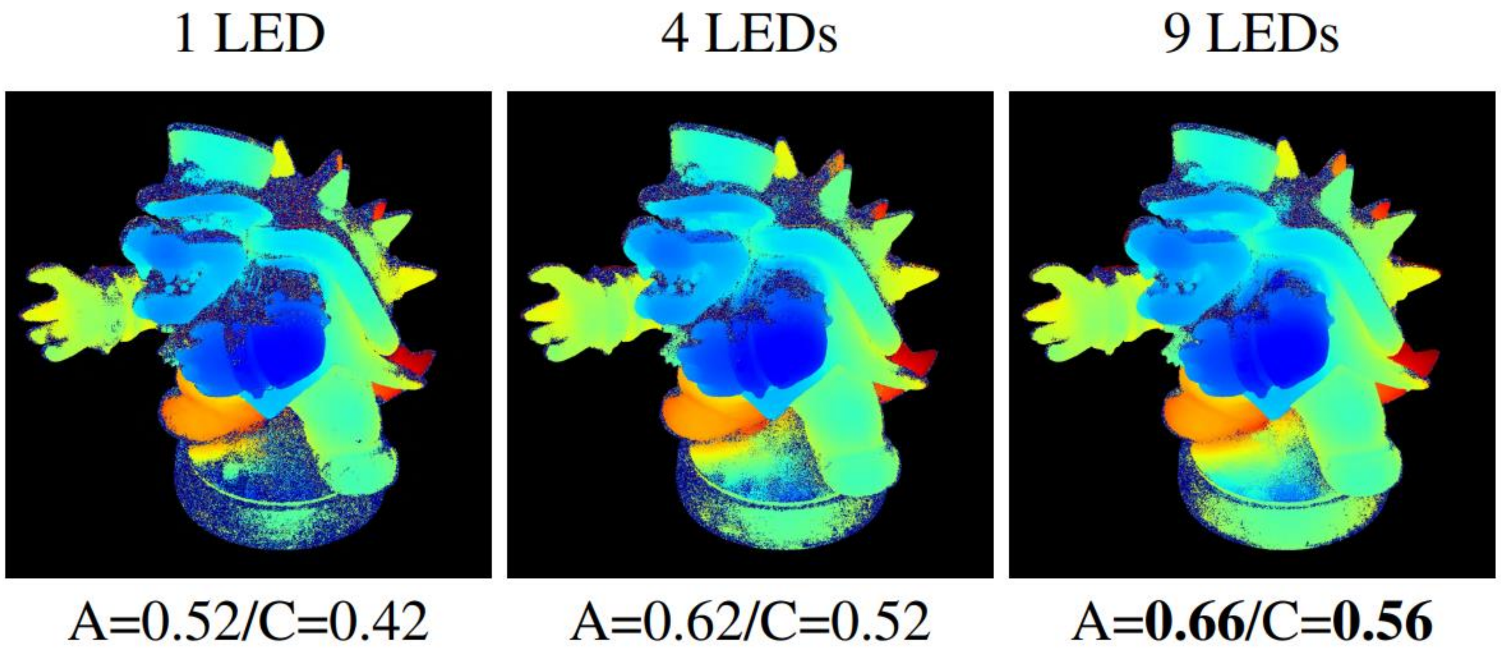

6 Results

Ablation Study

- Our learned patterns outperform the same number of randomly point lights

Photo | 32 Point Patterns | 24 Optimized Patterns | 32 Optimized Patterns | 48 Optimized Patterns

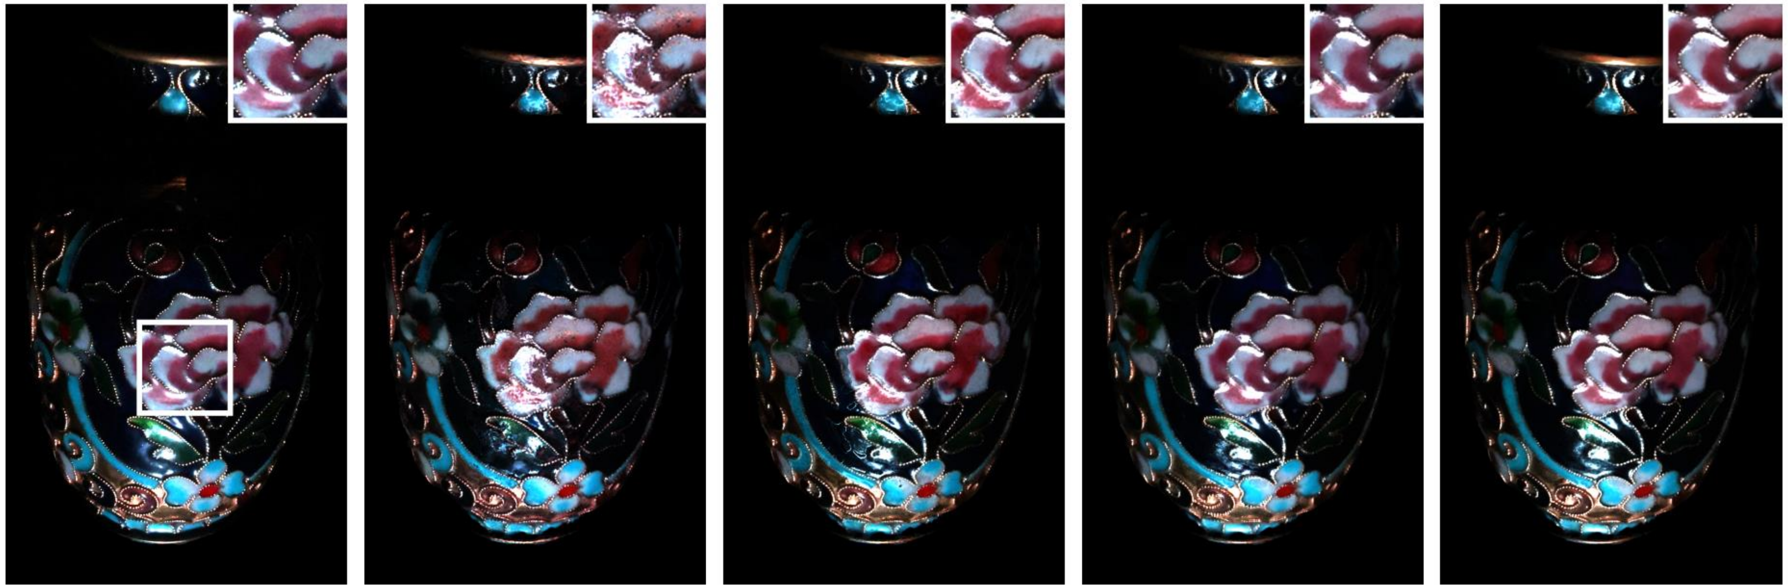

SSIM=0.95 | SSIM=0.95 | SSIM=0.96 | SSIM=0.96

Conclusion

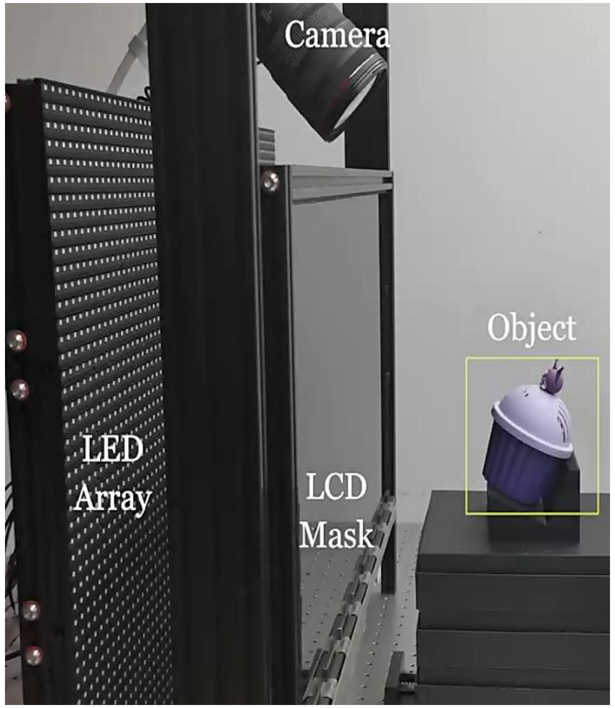

Acquisition Setup &
Differentiable Calibration

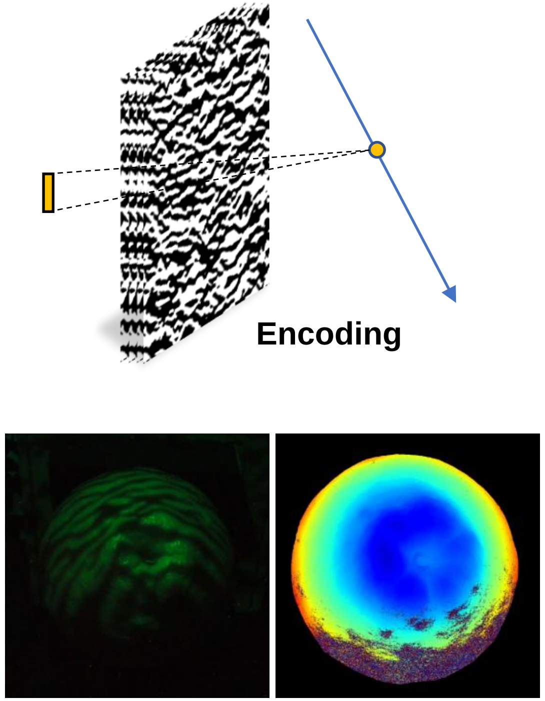

Structure Light with
Area Light

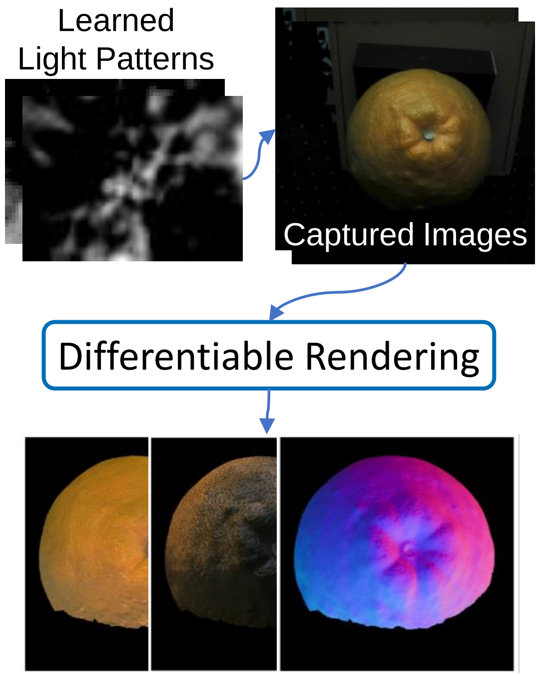

Reflectance Fine-tuning

7 Limitations and Future Work

Limitations

- Depth acquisition requires a long exposure time
- Multiple LEDs are manually selected
- Inter-reflection is not considered

Future Work

- Joint multiplexing of both LEDs and mask
- Adaptive acquisition pipeline
- Develop a handheld scanner with lightfield structured light

Thanks