Neural Rendering

CS 482 Interactive Computer Graphics

Jaeyoon Kim (TA)

KAIST

Contents

- **Topics in Neural Rendering**
- **NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis**

Neural Rendering

- **Deep Learning + Rendering = Neural Rendering**

- **A very new and rapidly emerging field**

Topics in Neural Rendering

- **Semantic Photo Synthesis and Manipulation**

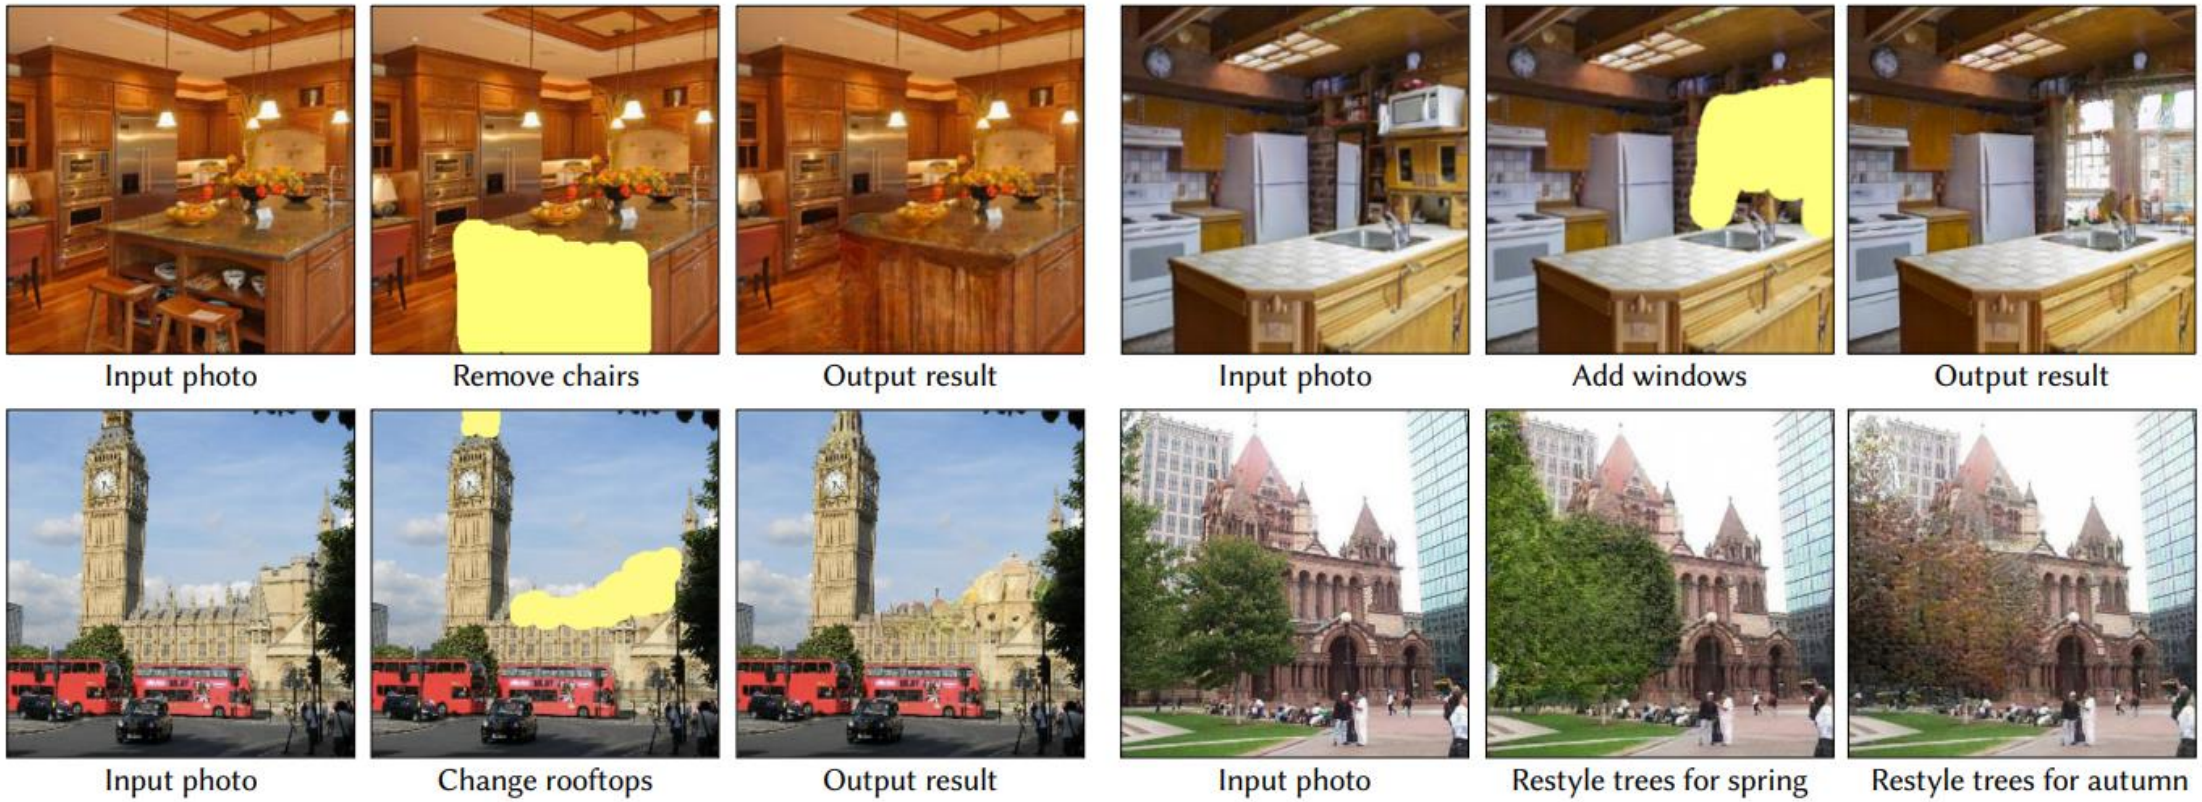

Topics in Neural Rendering

- **Semantic Photo Synthesis and Manipulation**
- **View Synthesis**

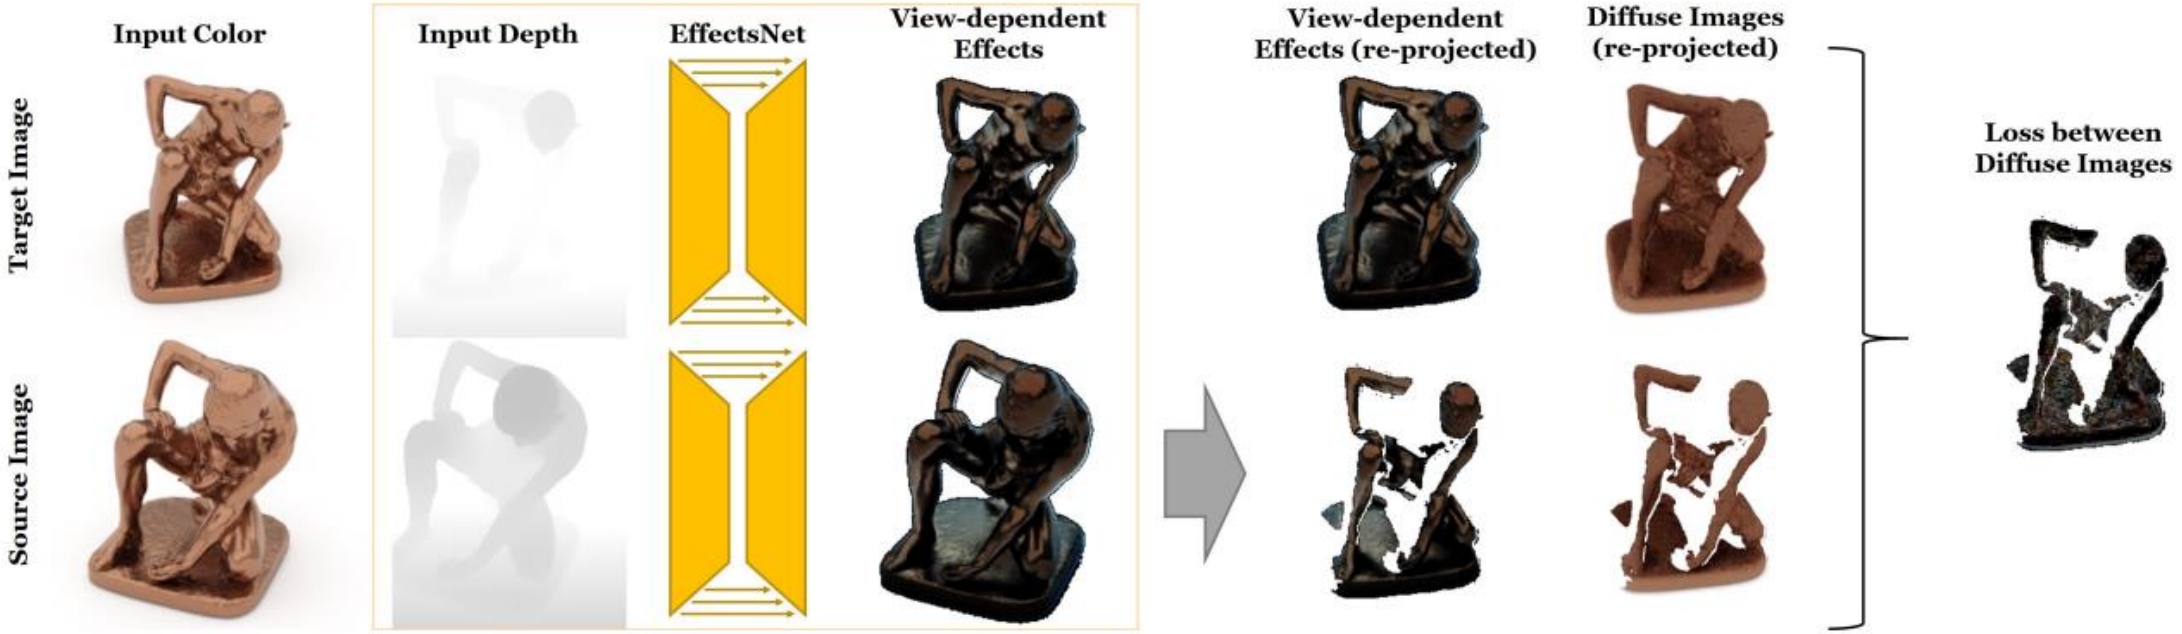

Topics in Neural Rendering

- **Semantic Photo Synthesis and Manipulation**
- **View Synthesis**
- **Relighting**

Topics in Neural Rendering

- **Semantic Photo Synthesis and Manipulation**
- **View Synthesis**
- **Relighting**
- **Facial/Body Reenactment**

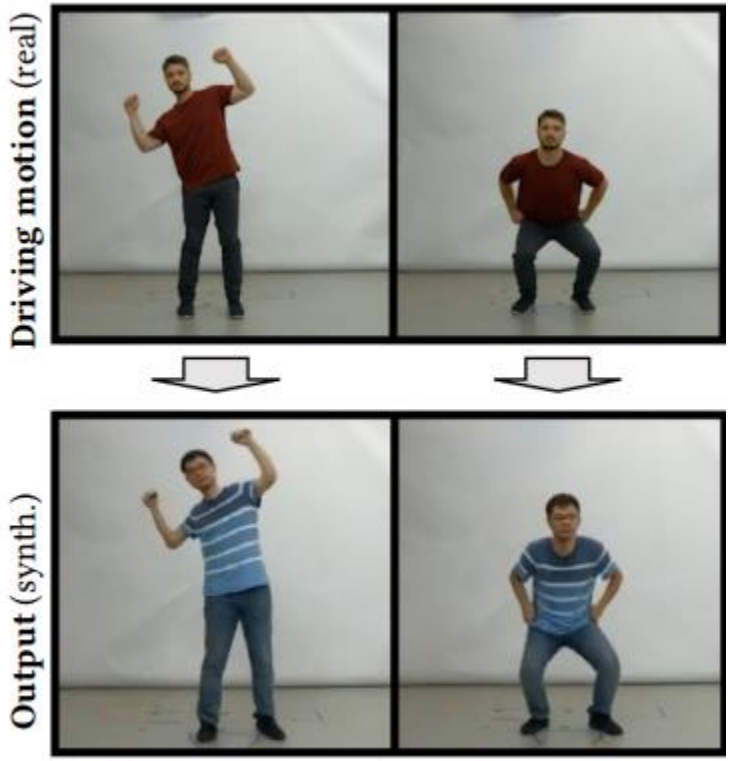

KAIST

Topics in Neural Rendering

- **Semantic Photo Synthesis and Manipulation**
- **View Synthesis**
- **Relighting**
- **Facial/Body Reenactment**
- **Anything else combining deep learning and CG concepts!**

NeRF – Neural Radiance Fields

- A novel view image generation by neural radiance fields and volume rendering
- A fully-connected network representing radiance fields
- Trained from a set of images with known camera poses

KAIST

NeRF – Neural Radiance Fields

- **Radiance fields**
  - **A vector function representing radiance in every direction through every point in space**
  - $\boldsymbol{L}: \mathbb{R}^3 \times \boldsymbol{S}^2 \rightarrow \mathbb{R}^3$
  - $(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z}) \times (\boldsymbol{\theta}, \boldsymbol{\phi}) \mapsto (\boldsymbol{r}, \boldsymbol{g}, \boldsymbol{b})$

NeRF – Neural Radiance Fields

- **Volume rendering**

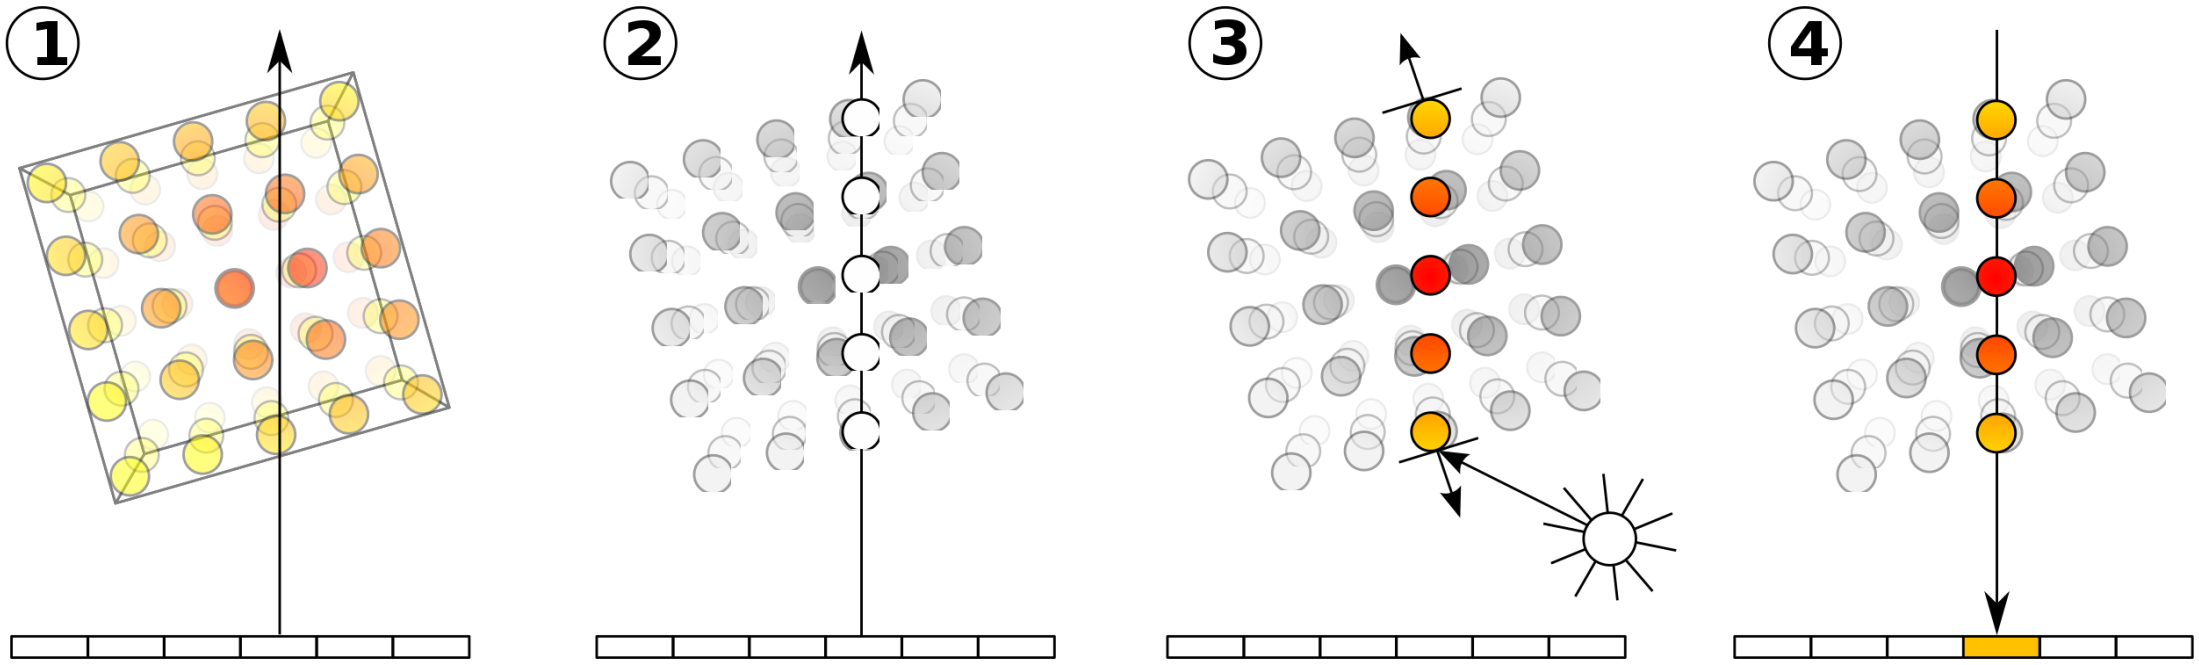

KAIST

NeRF – Neural Radiance Fields

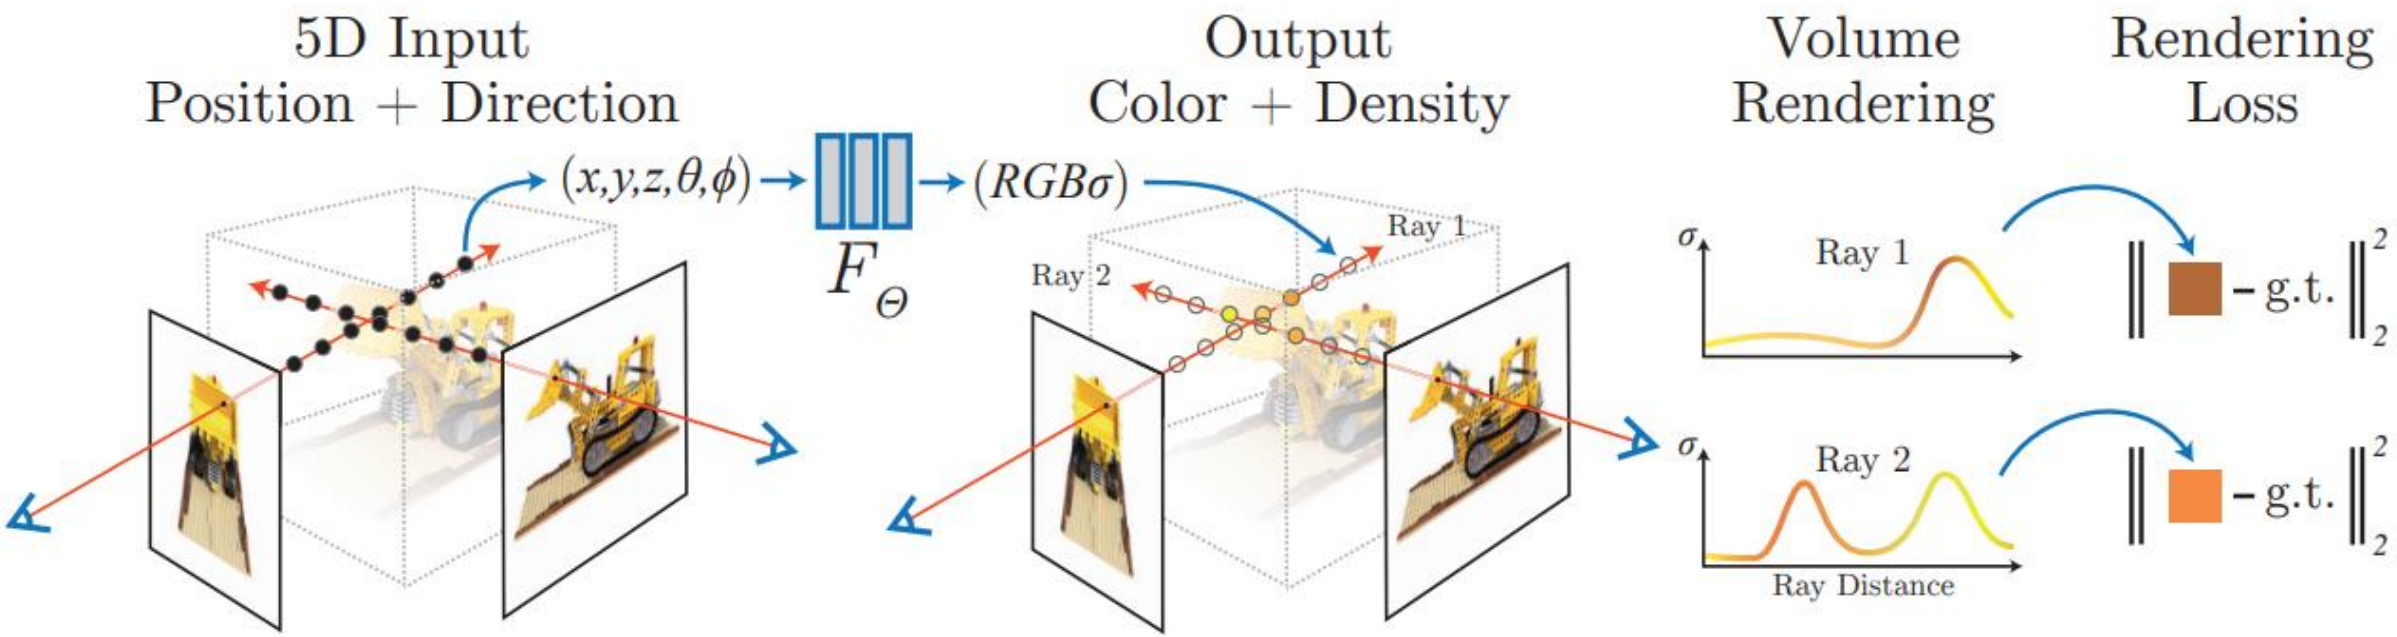

NeRF – Neural Radiance Fields

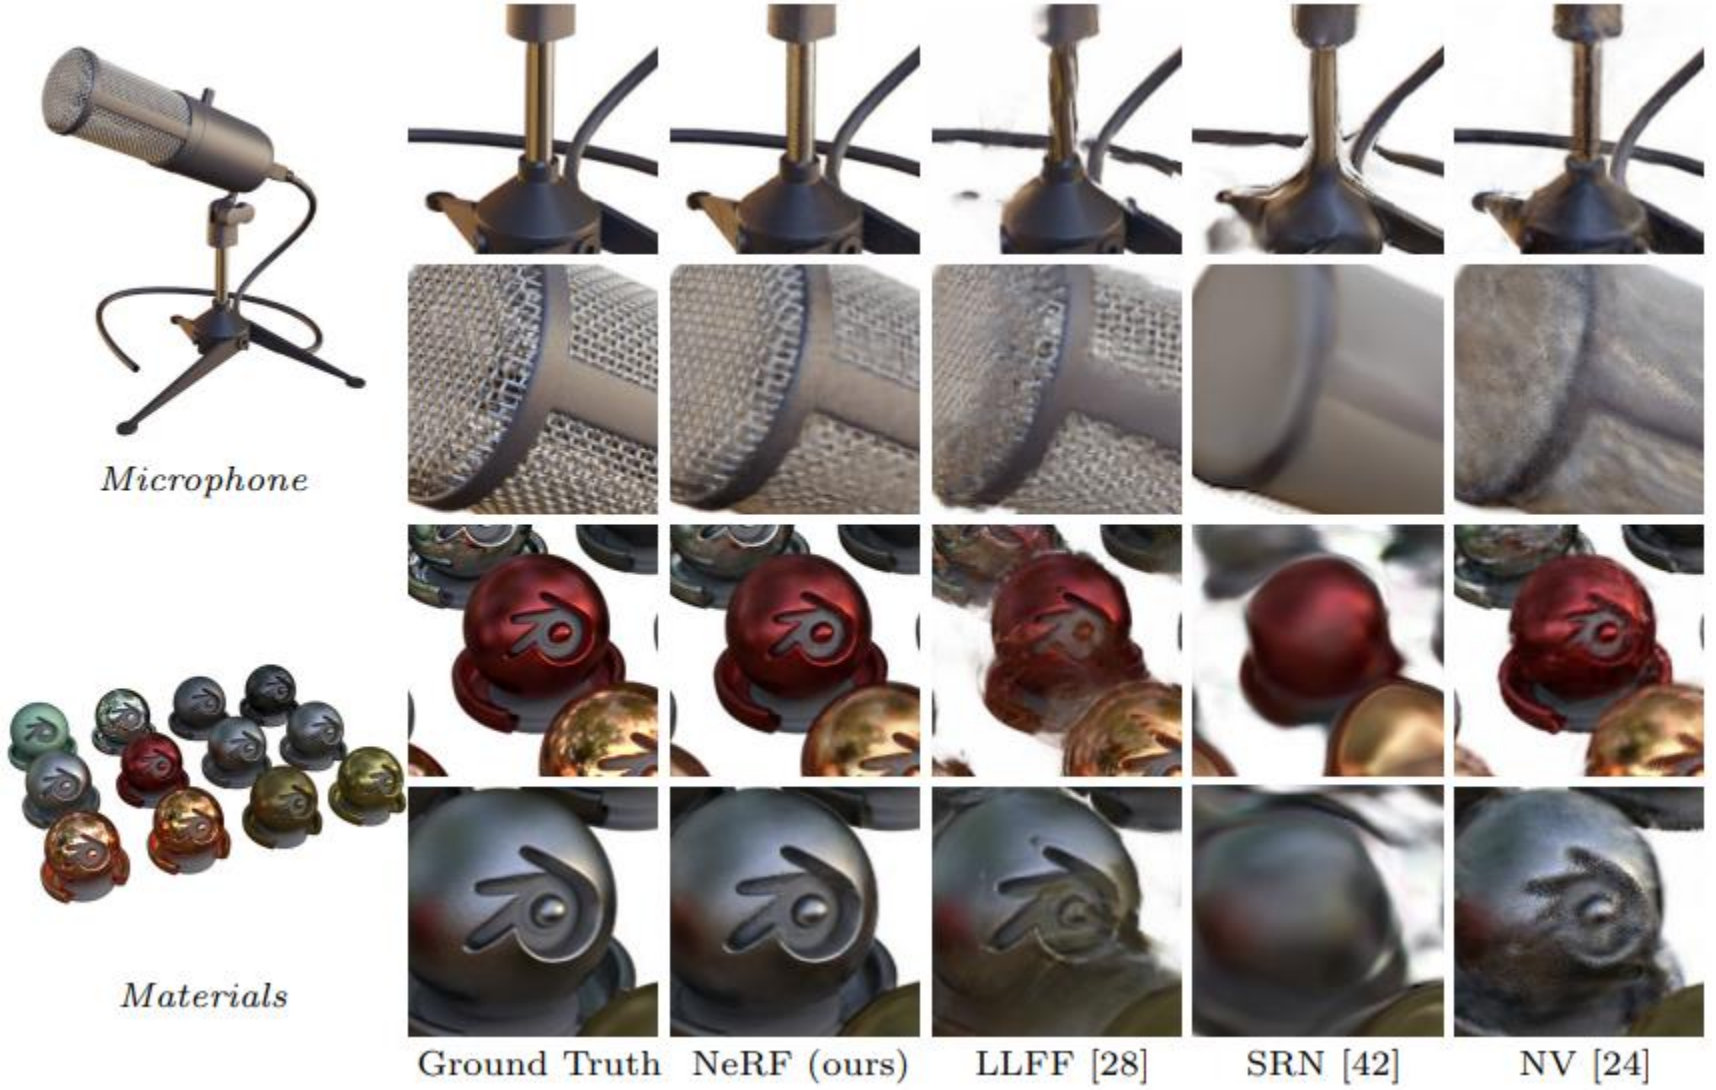

NeRF – Pros

- **High quality of details**
- **Low memory (~5MB for MLP Weights)**
  - **No other data required (e.g. Mesh, Voxel grids, etc.)**

Slide credit from Inyoung Cho, Paper presentation of NeRF from CS492 (Prof. Sung)

NeRF – Cons

- Slow training & inference
  - Training: 1-2 days for NVIDIA V100 GPU
  - Inference: few seconds for FHD image
- Train one model for one scene
  - No generalization
- Scalability to large scenes

Slide credit from Inyoung Cho, Paper presentation of NeRF from CS492 (Prof. Sung)

Further Discussion

- **Faster training/inference**
- **Generalization on various scenes**
- **Multiscale representation**
- **Better detail reconstruction via position encoding**
- **Extend task**

Slide credit from Kyubeom Han, Paper presentation of NeRF

Faster Training & Inference

- **DONeRF: Towards Real-Time Rendering of Compact Neural Radiance Fields using Depth Oracle Networks**
  - **Reduce required samples via predicting locations of rays through depths**
  - **48x inference cost, same quality with 4 samples compared to NeRF(256 samples)**

Slide credit from Kyubeom Han, Paper presentation of NeRF

Faster Training & Inference

- **DONeRF: Towards Real-Time Rendering of Compact Neural Radiance Fields using Depth Oracle Networks**

Slide credit from Kyubeom Han, Paper presentation of NeRF

Generalization on Scenes

- **Learned Initializations for Optimizing Coordinate-Based Neural Representations**
  - **Better initialization via meta-learning**

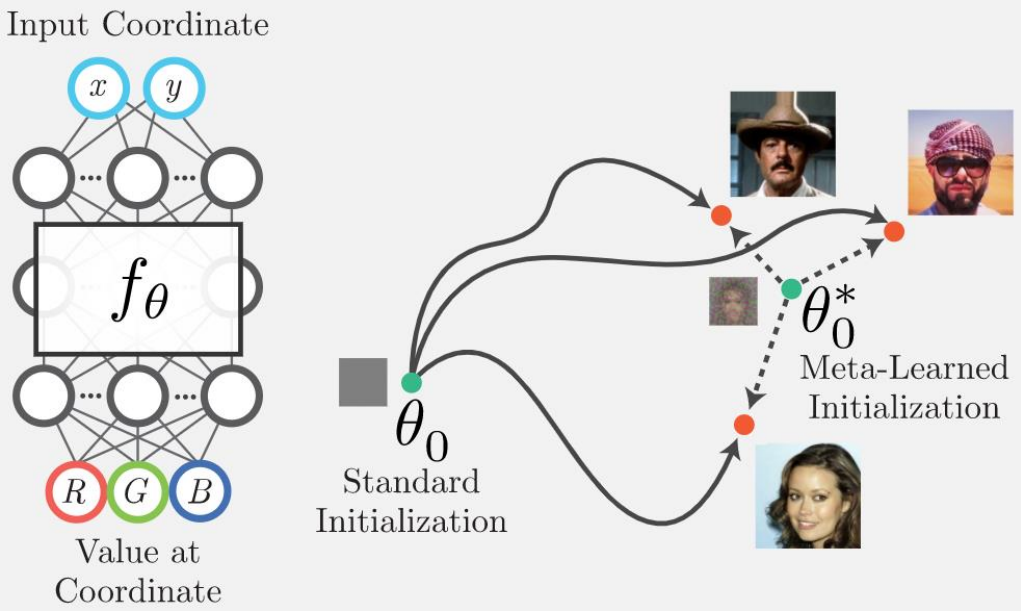

Slide credit from Kyubeom Han, Paper presentation of NeRF

Multiscale Representation

- **Mip-NeRF: A Multiscale Representation for Anti-Aliasing Neural Radiance Fields**
  - **Mipmap for various scale representation and better positional encoding**

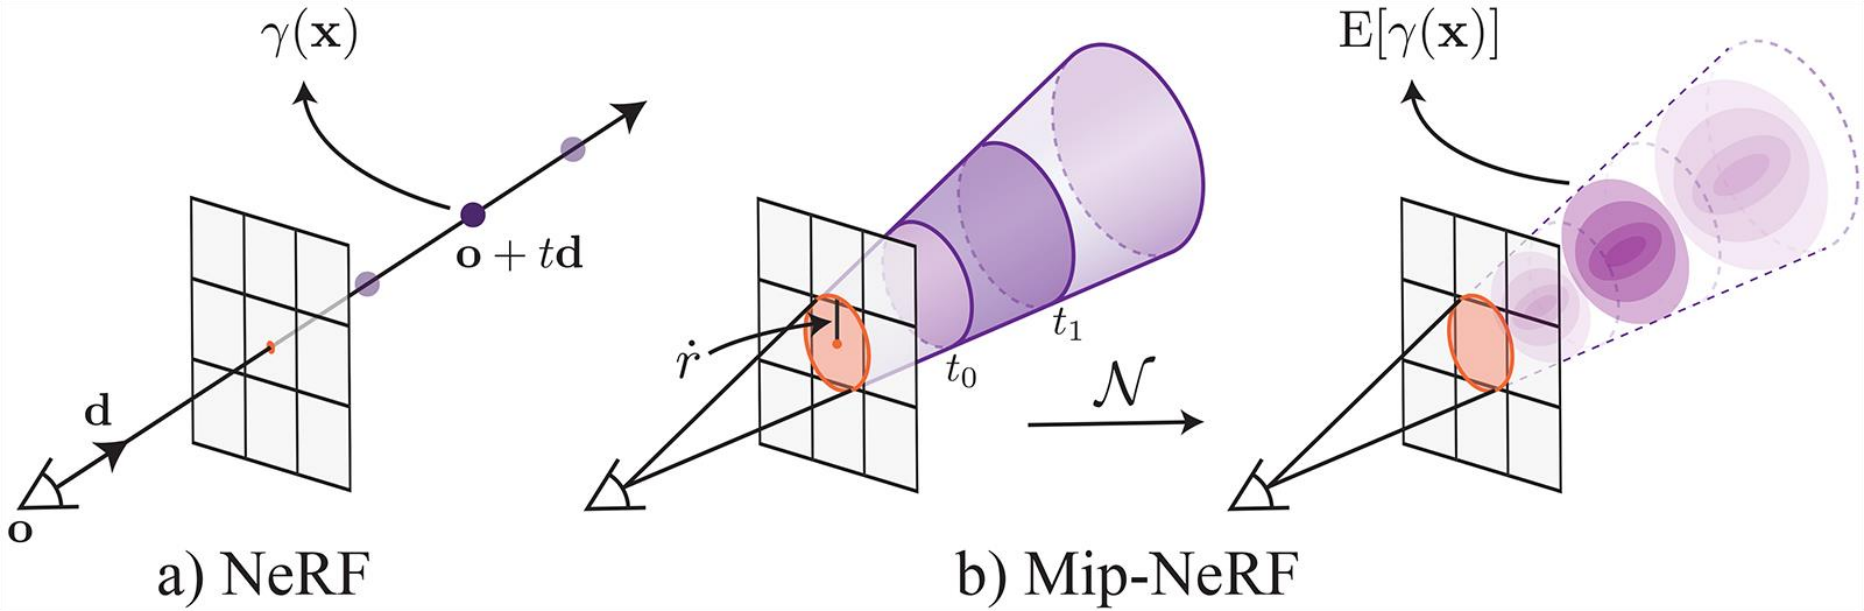

Slide credit from Kyubeom Han, Paper presentation of NeRF

KAIST

Multiscale Representation

- **Mip-NeRF: A Multiscale Representation for Anti-Aliasing Neural Radiance Fields**

Slide credit from Kyubeom Han, Paper presentation of NeRF

NeRF for Relighting

- **NeRV: Neural Reflectance and Visibility Fields for Relighting and View Synthesis**
  - **Predicting reflectance & visibility field to deal with arbitrary lighting**

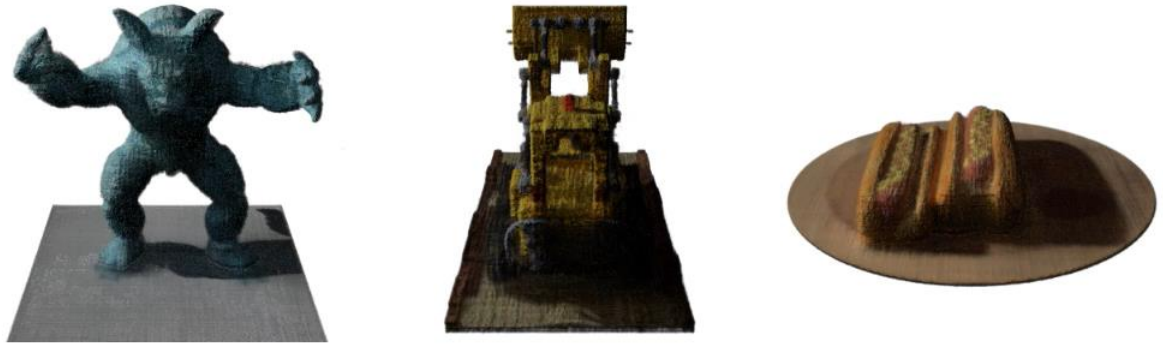

Slide credit from Kyubeom Han, Paper presentation of NeRF

Summary

- **Neural rendering**

- **Recent neural volume rendering techniques for view synthesis**

Reference

- Tewari et al. "State of the art on neural rendering." Computer Graphics Forum. Vol. 39. No. 2. 2020.
- Bau et al. "Semantic photo manipulation with a generative image prior." ACM Trans. Graph. 38, 4 (2019).
- Thies et al. “Image-guided neural object rendering.” ICLR 2020.
- Meka et al. “Deep reflectance fields: High-quality facial reflectance field inference from color gradient illumination.” ACM Trans. Graph. 38, 4 (2019).
- Liu et al. “Neural rendering and reenactment of human actor videos.” ACM Trans. Graph. 38, 5 (2019).
- Fried et al. “Text-based editing of talking-head video.” ACM Trans. Graph. 38, 4 (2019).
- Neff et al. “DONeRF: Towards Real-Time Rendering of Compact Neural Radiance Fields using Depth Oracle Networks.” EGST 2021.
- Barron et al. “Mip-NeRF: A Multiscale Representation for Anti-Aliasing Neural Radiance Fields.” ICCV 2021 (Oral)
- Srinivasan et al. “NeRV: Neural Reflectance and Visibility Fields for Relighting and View Synthesis.” CVPR 2021.

KAIST