Generalization of Neural Fields
(Overview)

Motivation: Novel View Synthesis

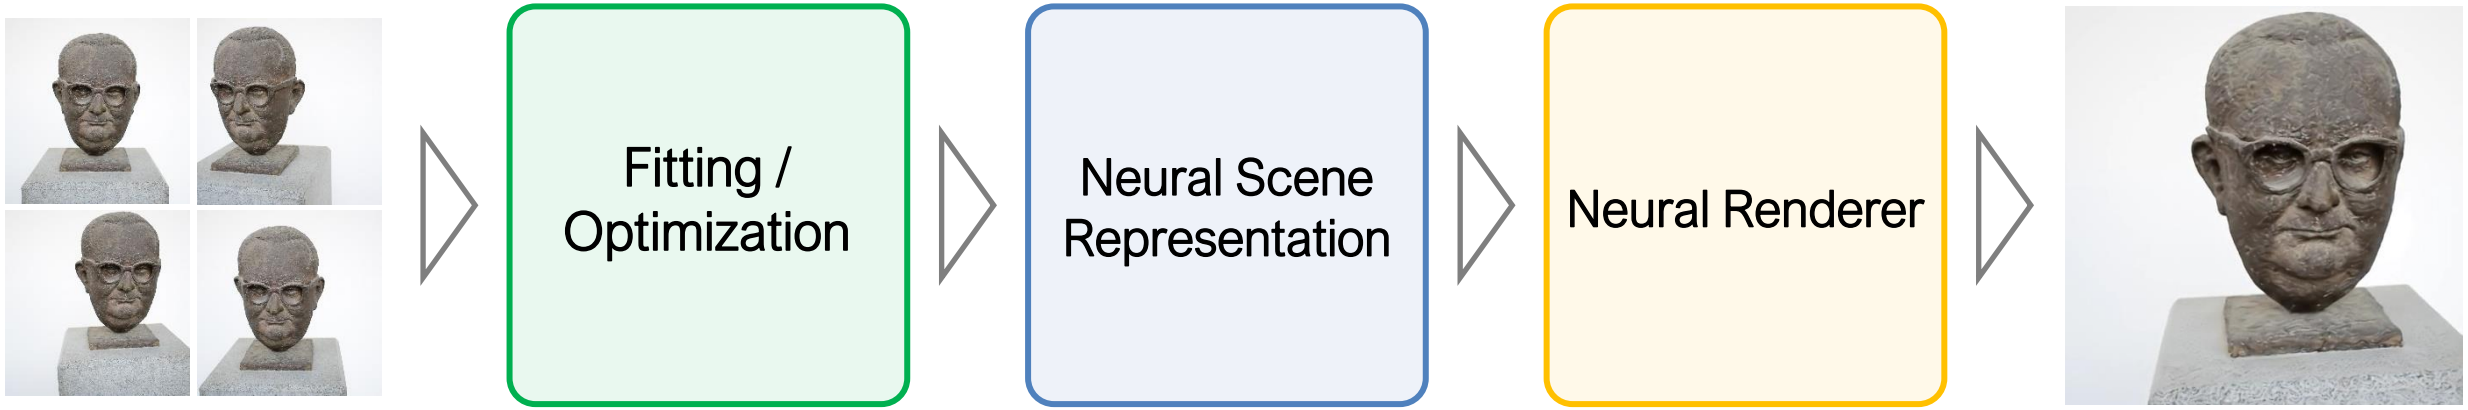

Motivation: Novel View Synthesis

DeepVoxels, CVPR 2018.

NeRF, ECCV 2021

IDR, ECCV 2021

Plenoxels, CVPR 2022

Inferring Neural Fields

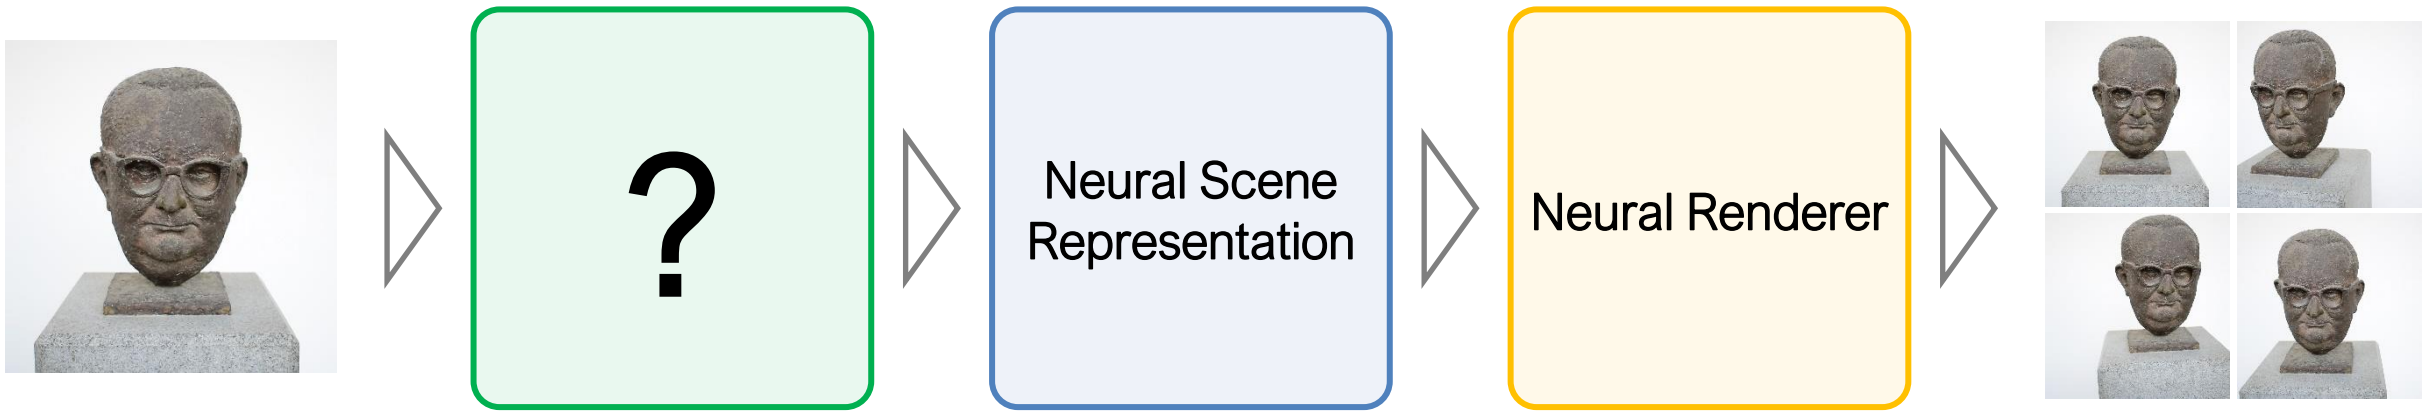

If only a single observation is available, or if only part of the scene has been observed, Inference needs to be prior-based – i.e., we need to learn to reconstruct.

General Framework: Encoder-Decoder

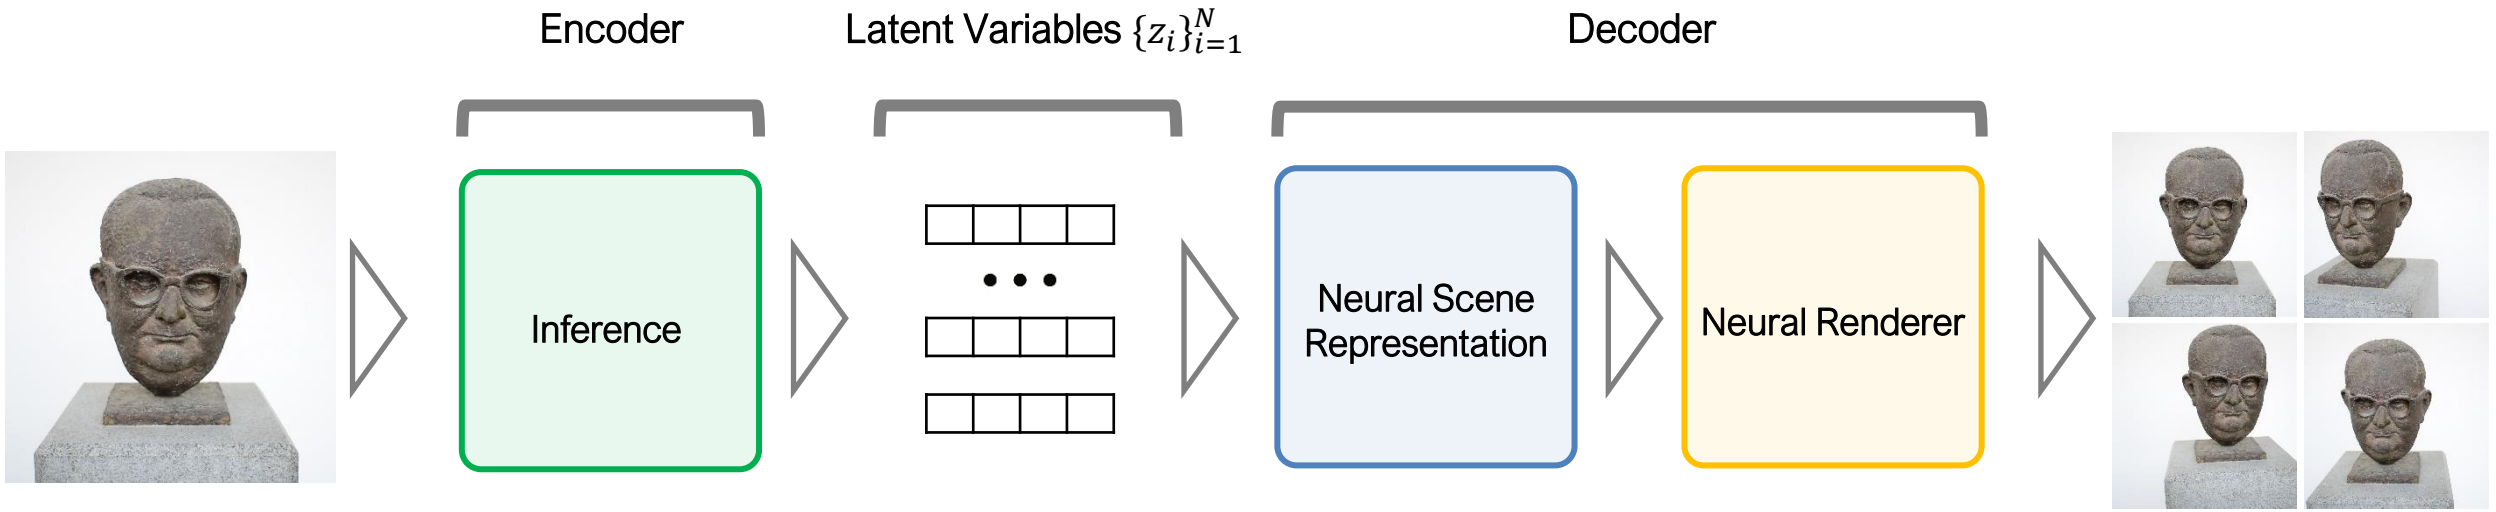

What are the latent variables?

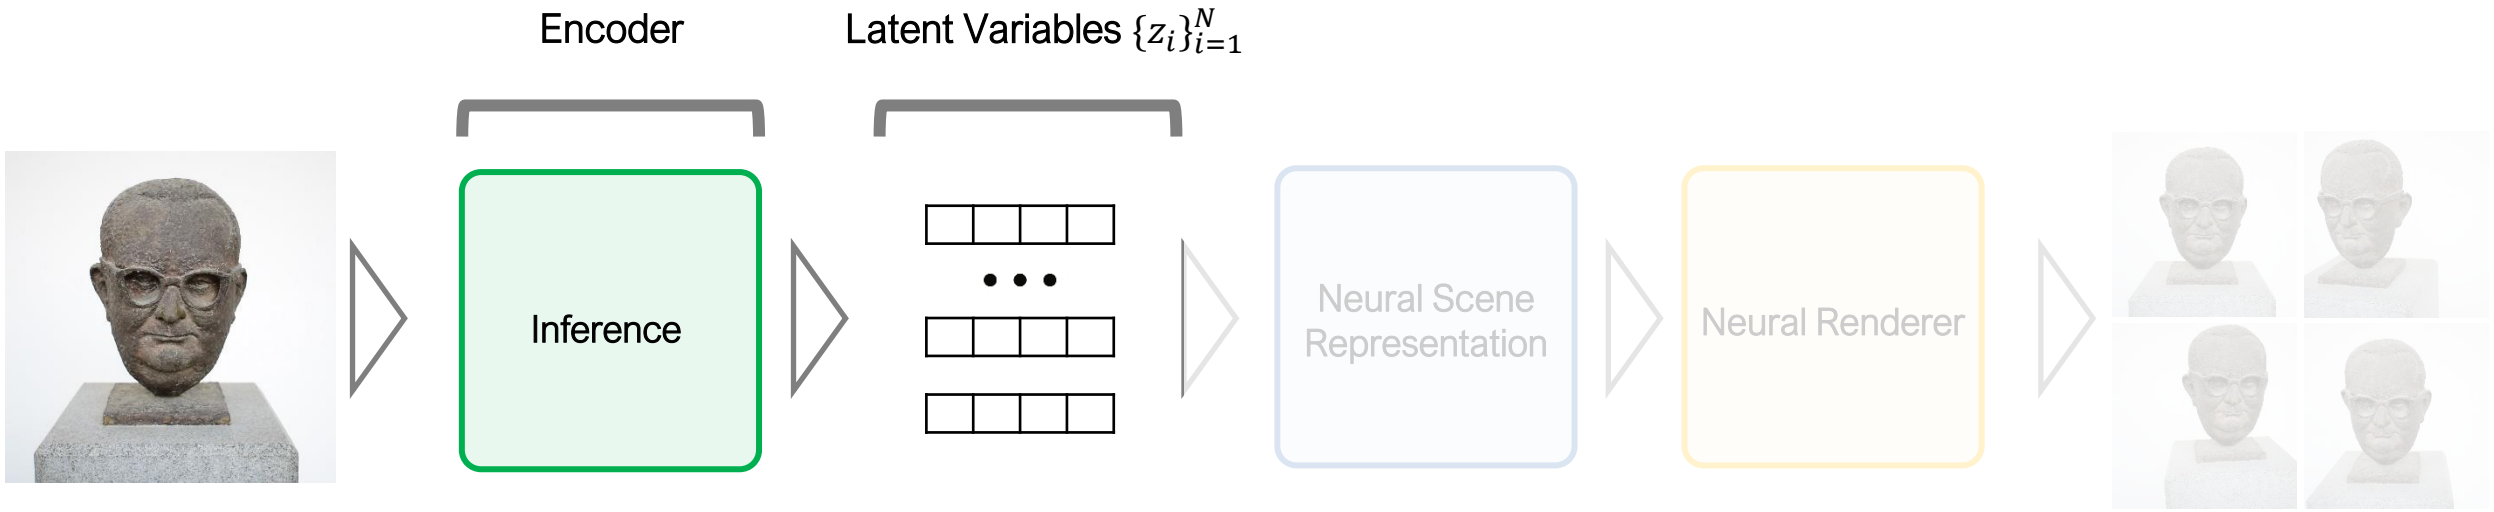

How to predict latent variables from observations?

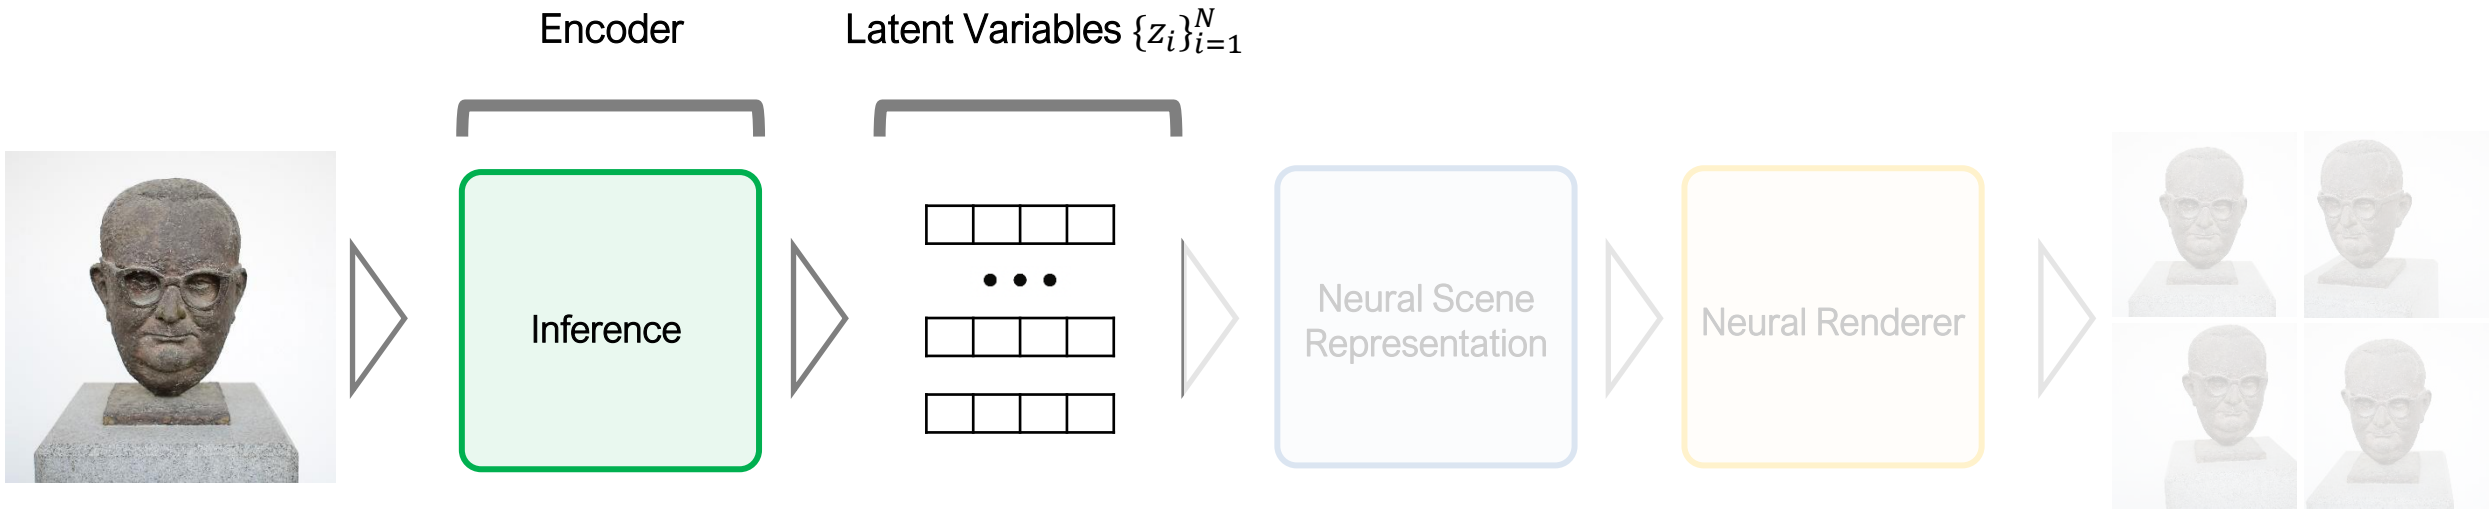

How do we decode latent variables into the Neural Field?

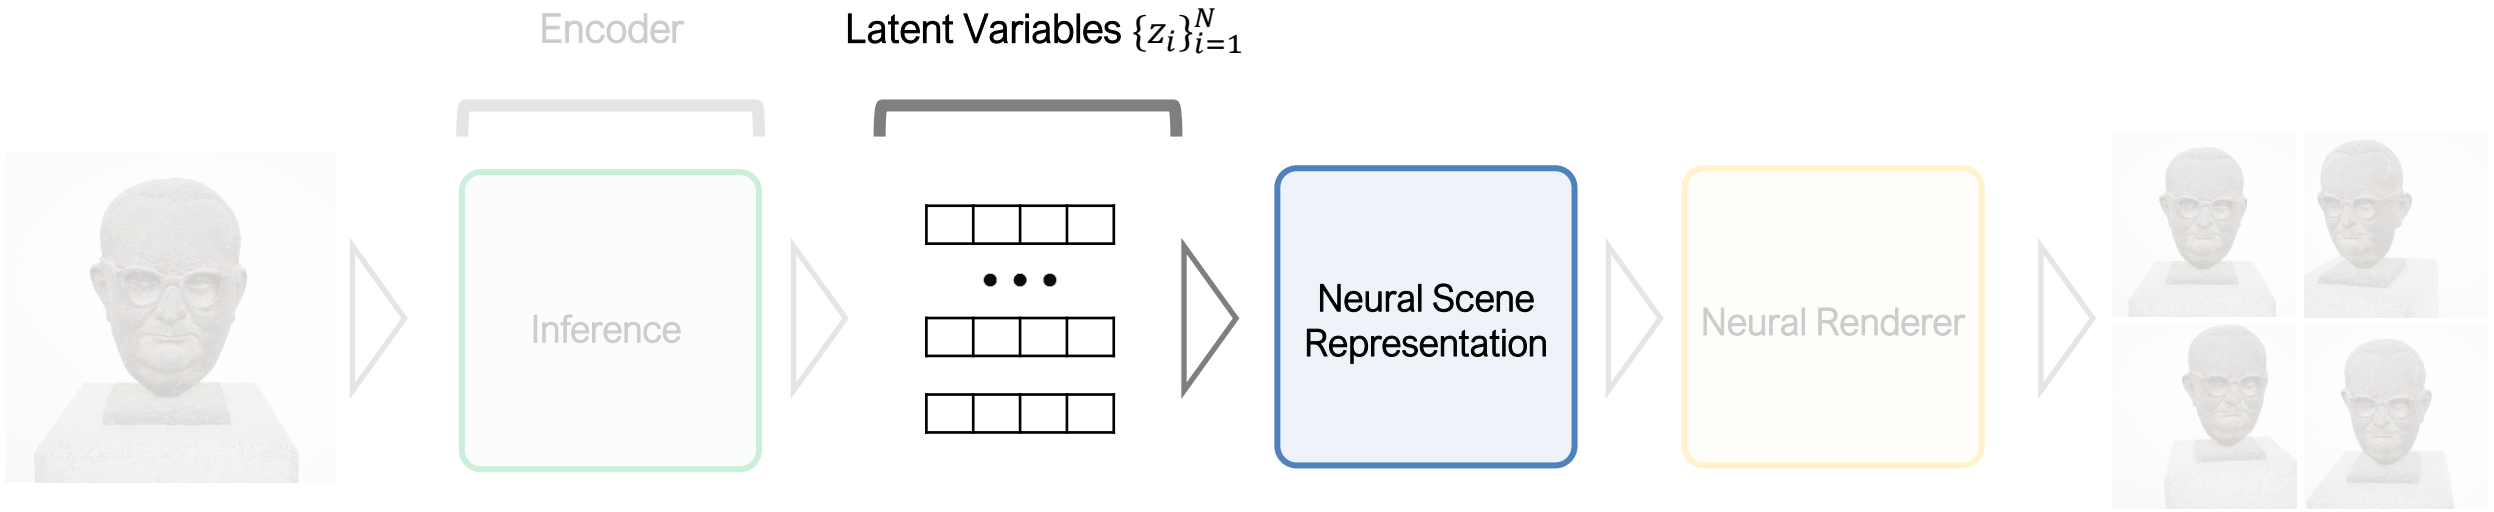

What are the latent variables?

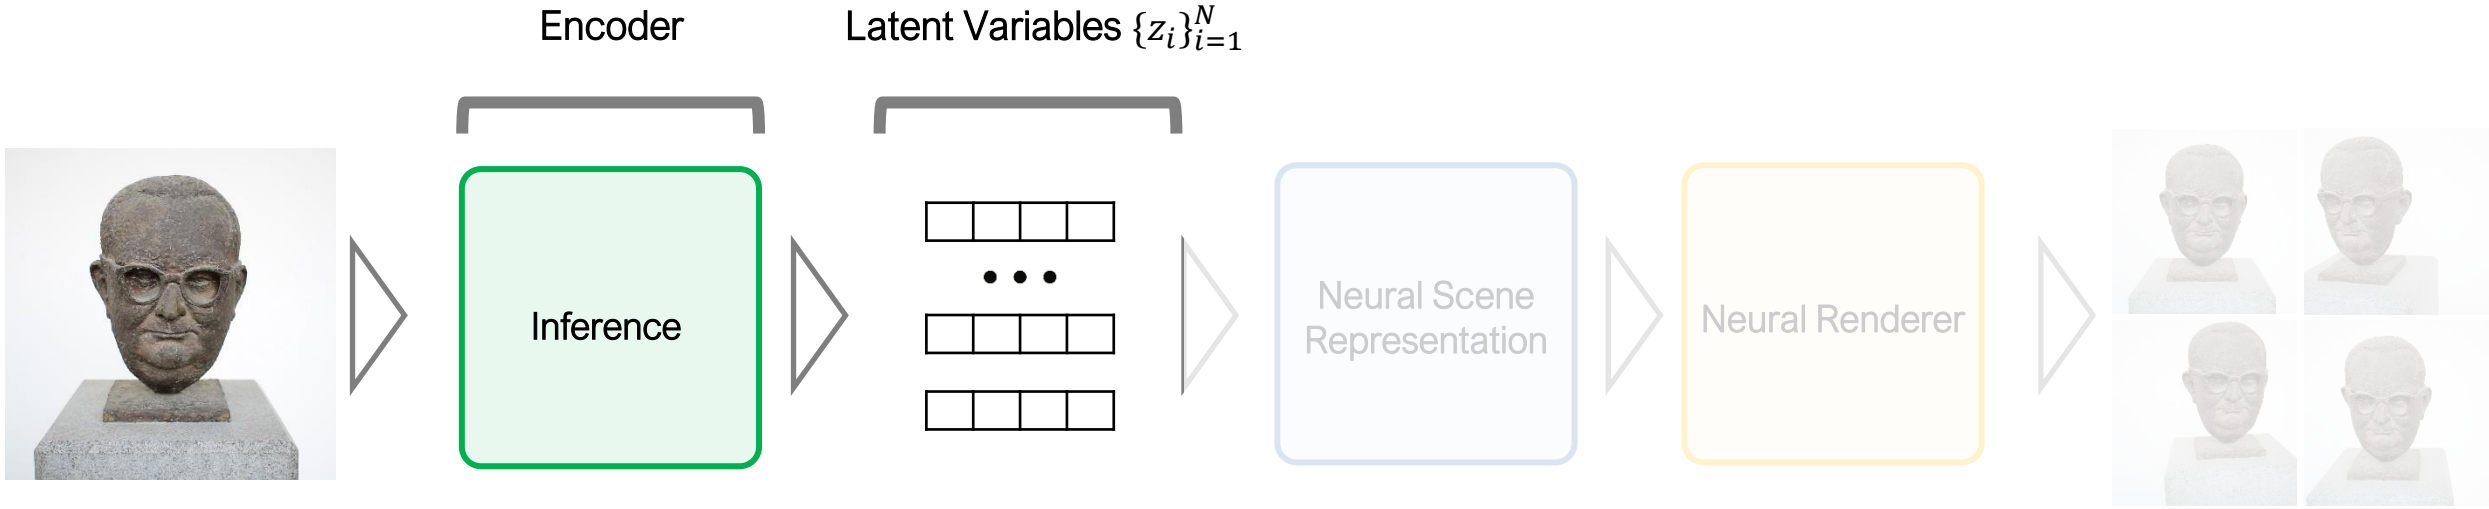

Key Consideration: *Locality.*

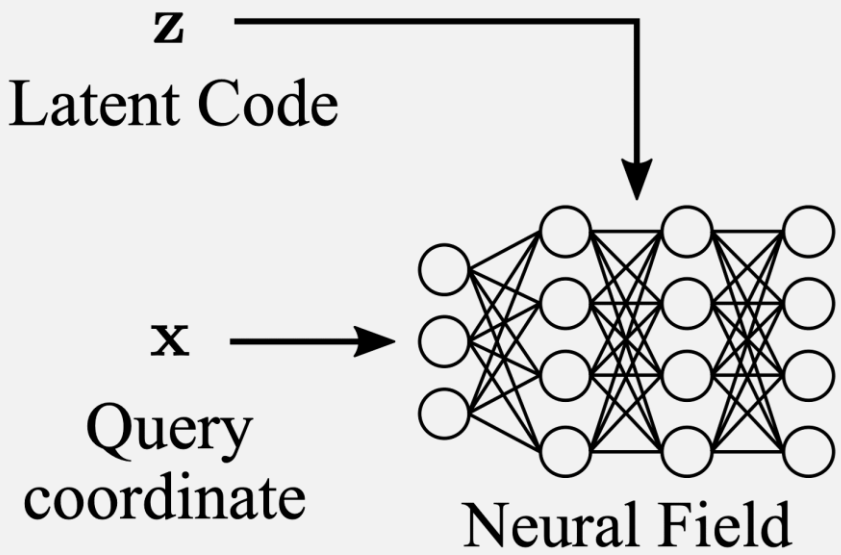

Global Conditioning

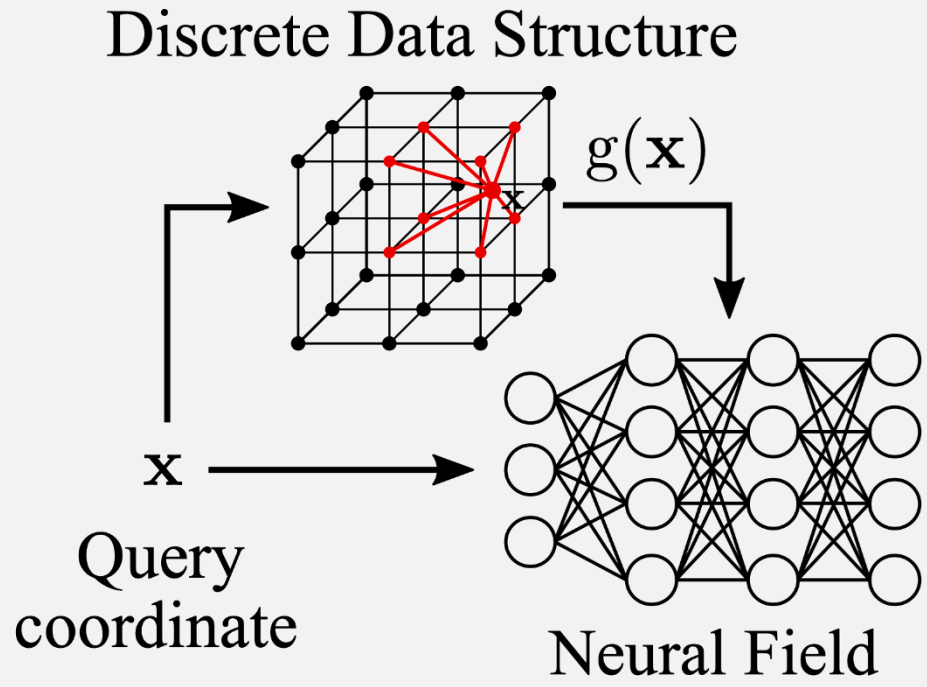

Local Conditioning

Neural Fields in Visual Computing and Beyond, Xie et al., EG STAR 2022

Global Latent Codes

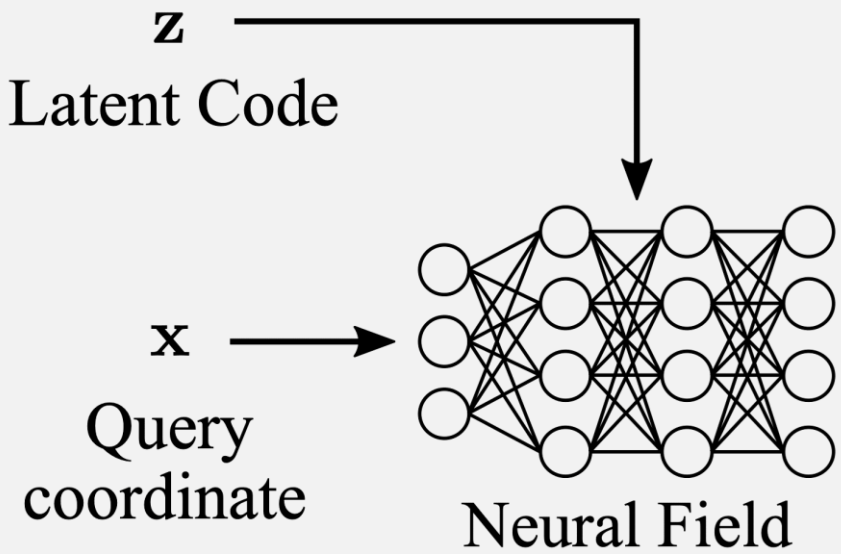

Global Conditioning

Local Conditioning

Neural Fields in Visual Computing and Beyond, Xie et al., EG STAR 2022

Global Conditioning

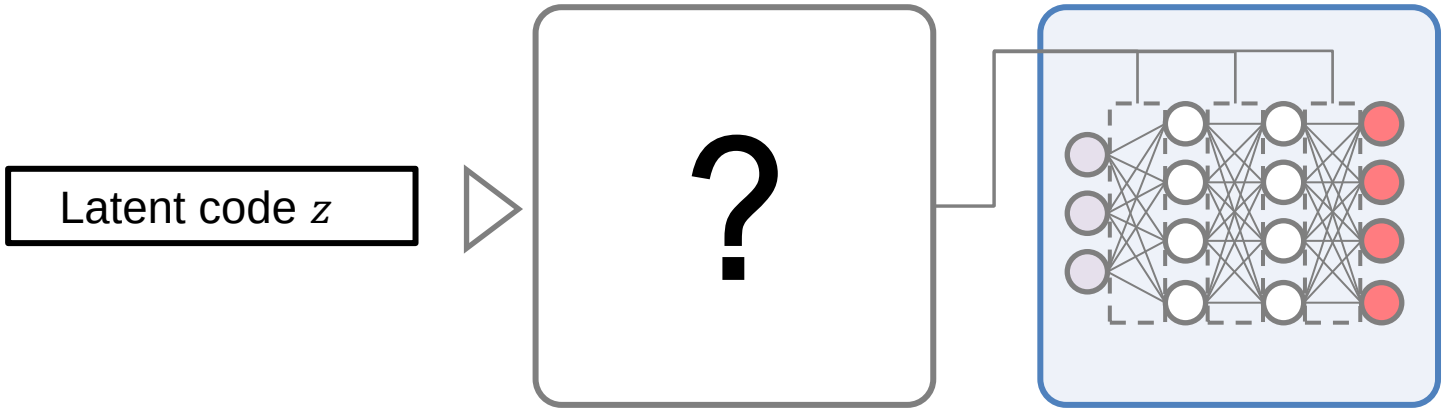

Global Conditioning

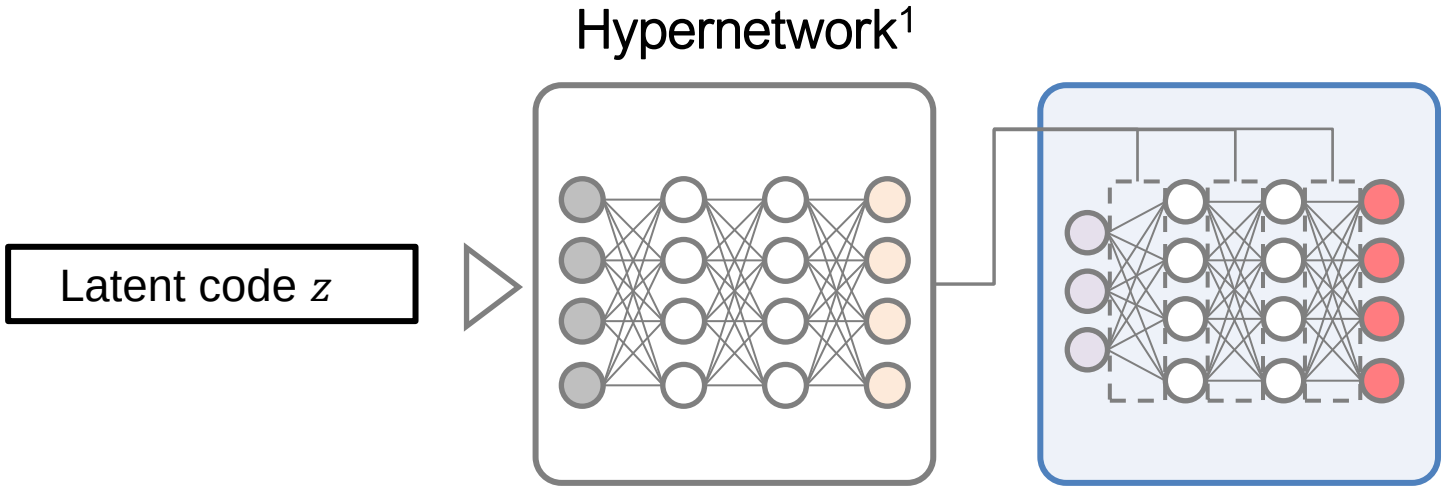

[1][Schmidhuber et al. 1992, Schmidhuber et al. 1993, Stanley et al. 2009, Ha et al., 2016]

Global Latent Codes:
Enables reconstruction from *partial* observations!

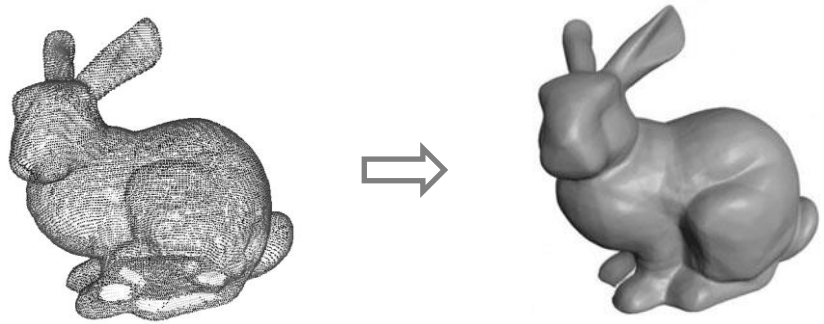

DeepSDF, Occupancy Networks, IM-Net

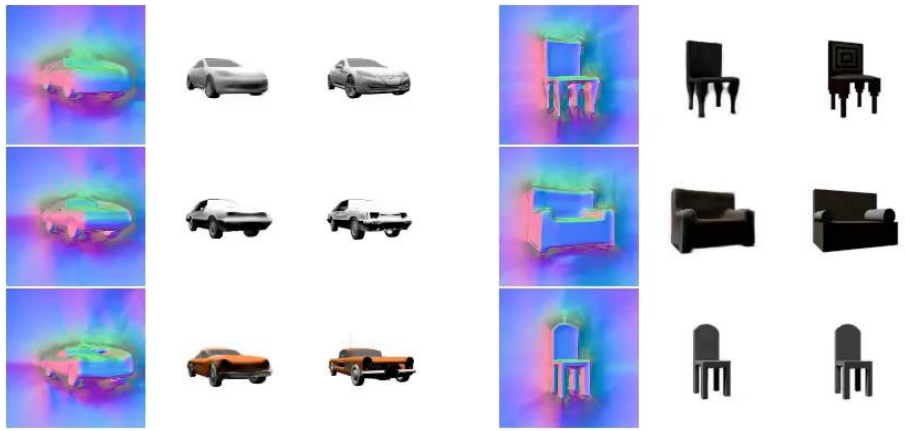

Scene Representation Networks: Continuous
3D-Structure-Aware Neural Scene Representations, NeurIPS 2019.

Differential Volumetric Rendering,
Niemeyer et al., CVPR 2020

Global Latent Codes:

Enables reconstruction from *partial* observations!

DeepSDF, Occupancy Networks, IM-Net

Key limitation: Simple, non-compositional scenes.
But: Latent Space for full objects (interpolation etc.)

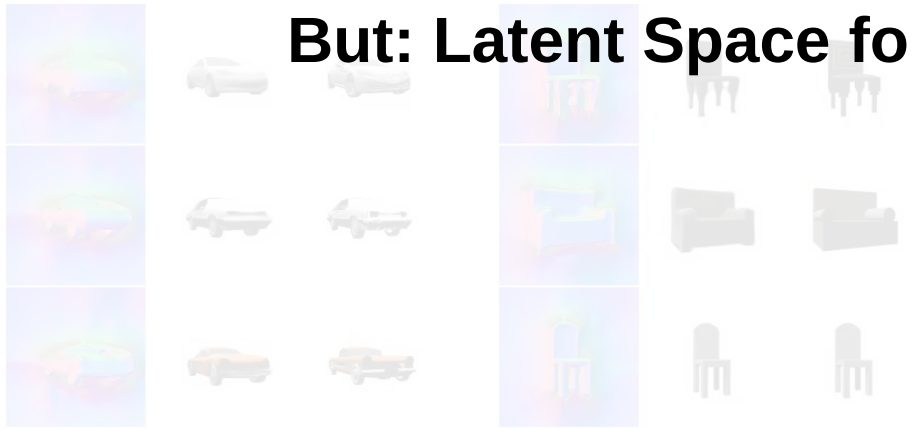

Scene Representation Networks: Continuous
3D-Structure-Aware Neural Scene Representations, NeurIPS 2019.

Differential Volumetric Rendering,
Niemeyer et al., CVPR 2020

Local Latent Codes

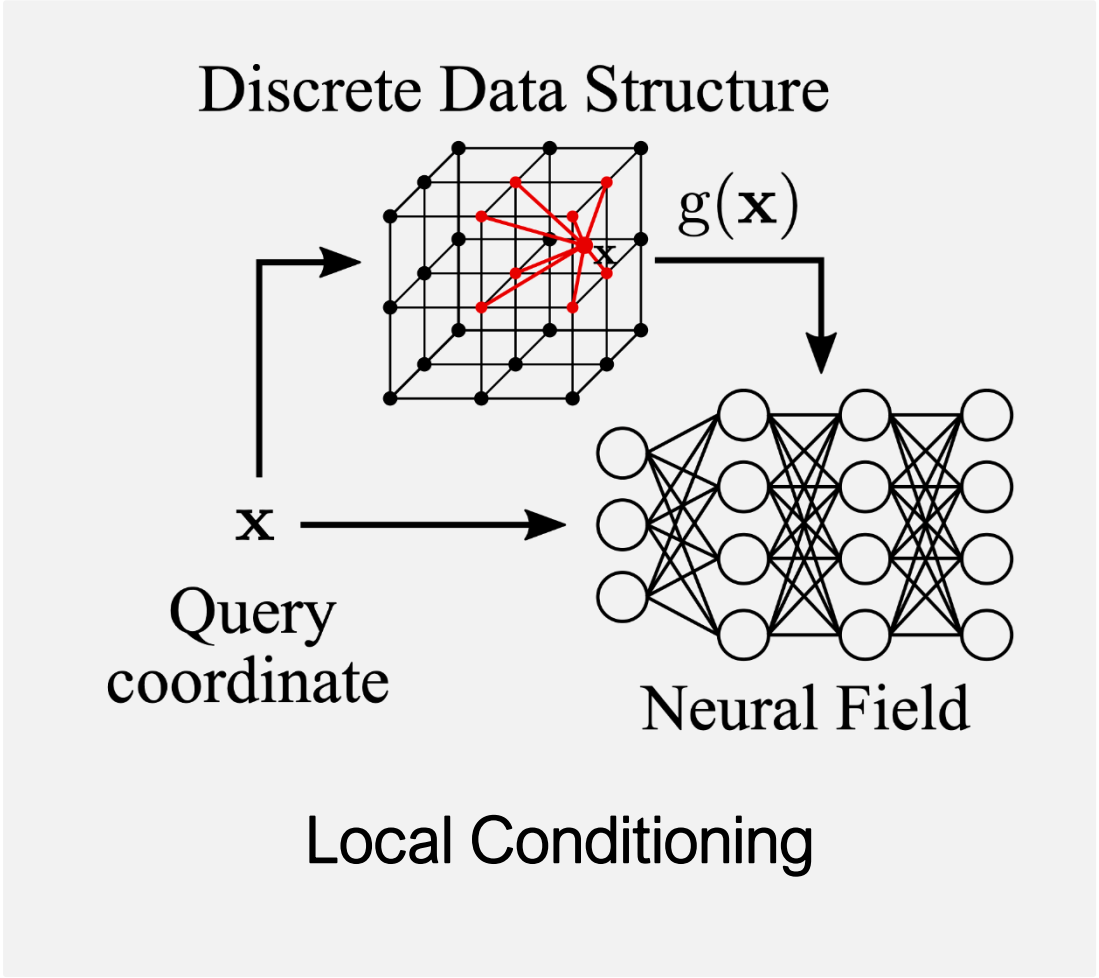

Neural Fields in Visual Computing and Beyond, Xie et al., EG STAR 2022

Local Conditioning: Pixel-Aligned Features

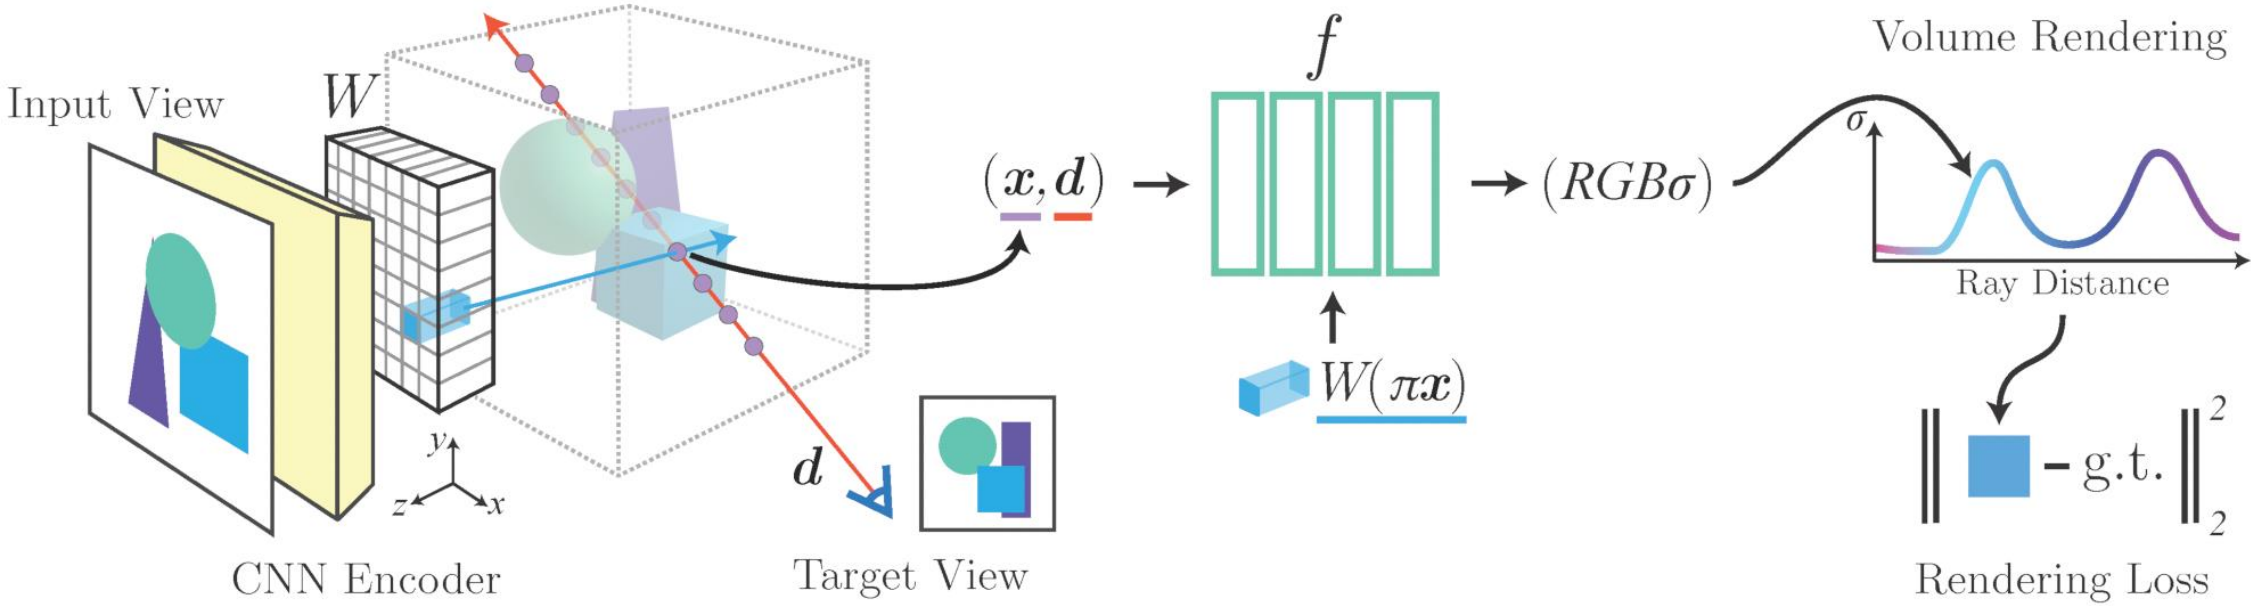

PiFU, Saito et al., ICCV 2019.
PixelNeRF, Yu et al., CVPR 2021
Grf: Learning a general radiance field…, Trevithick et al. ICCV 2021
Vision Transformer for NeRF-Based View Synthesis from a Single Input Image, Lin et al. 3DV 2022
MVSNeRF: Fast Generalizable Radiance Field Reconstruction from Multi-View Stereo, Chen et al. ICCV 2021
Neural Rays for Occlusion-aware Image-based Rendering, Liu et al. CVPR 2022
…

Local Conditioning: Pixel-Aligned Features

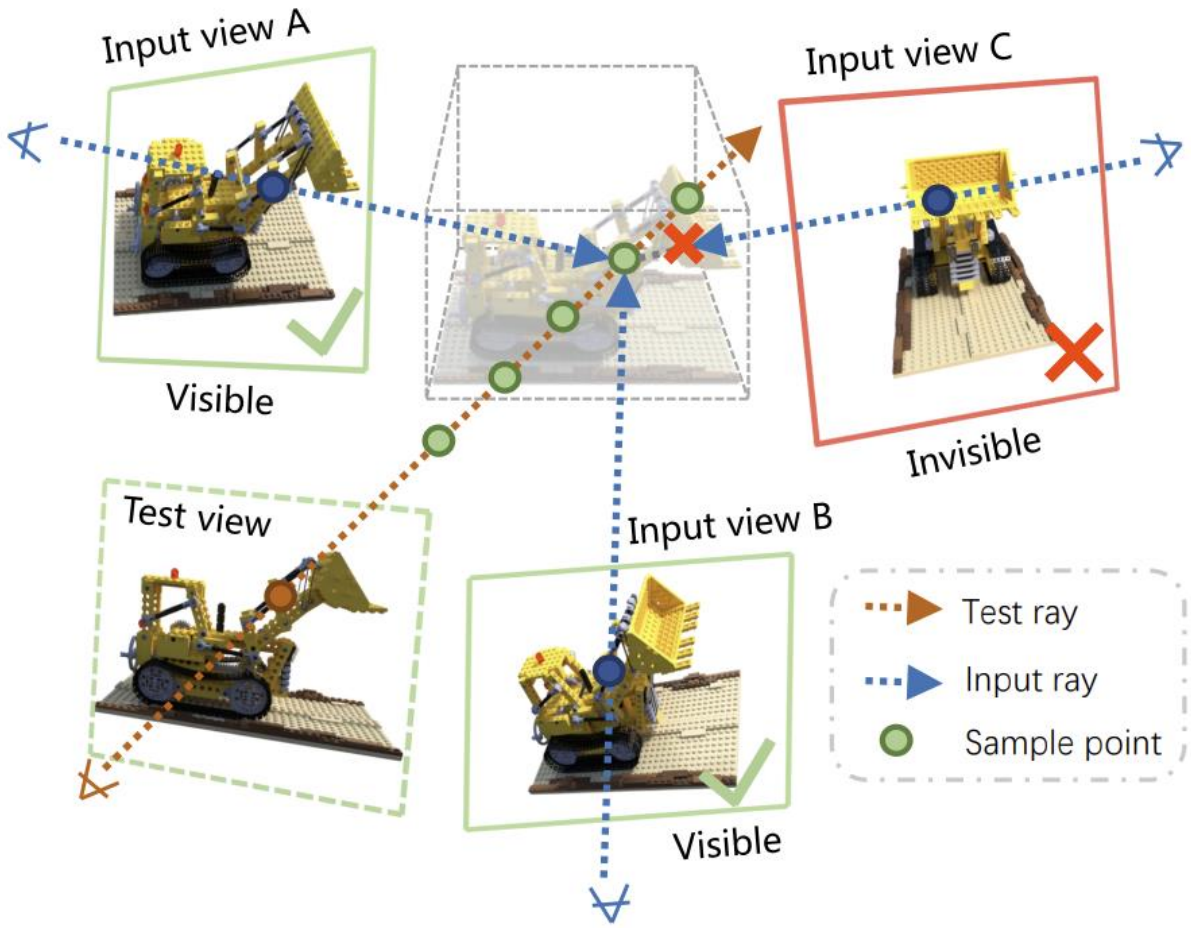

Neural Rays for Occlusion-aware Image-based Rendering, Liu et al. CVPR 2022

Local Conditioning: Pixel-Aligned Features

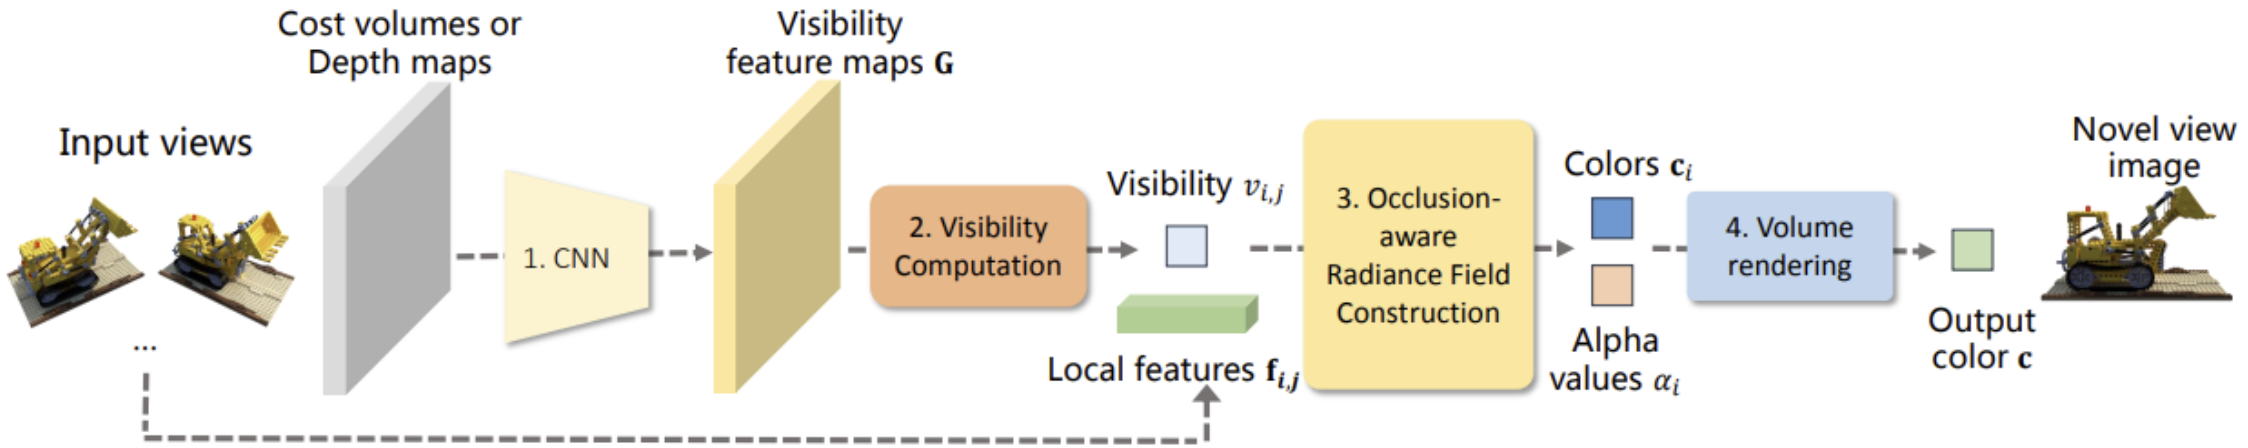

Neural Rays for Occlusion-aware Image-based Rendering, Liu et al. CVPR 2022

Local Conditioning: Pixel-Aligned Features

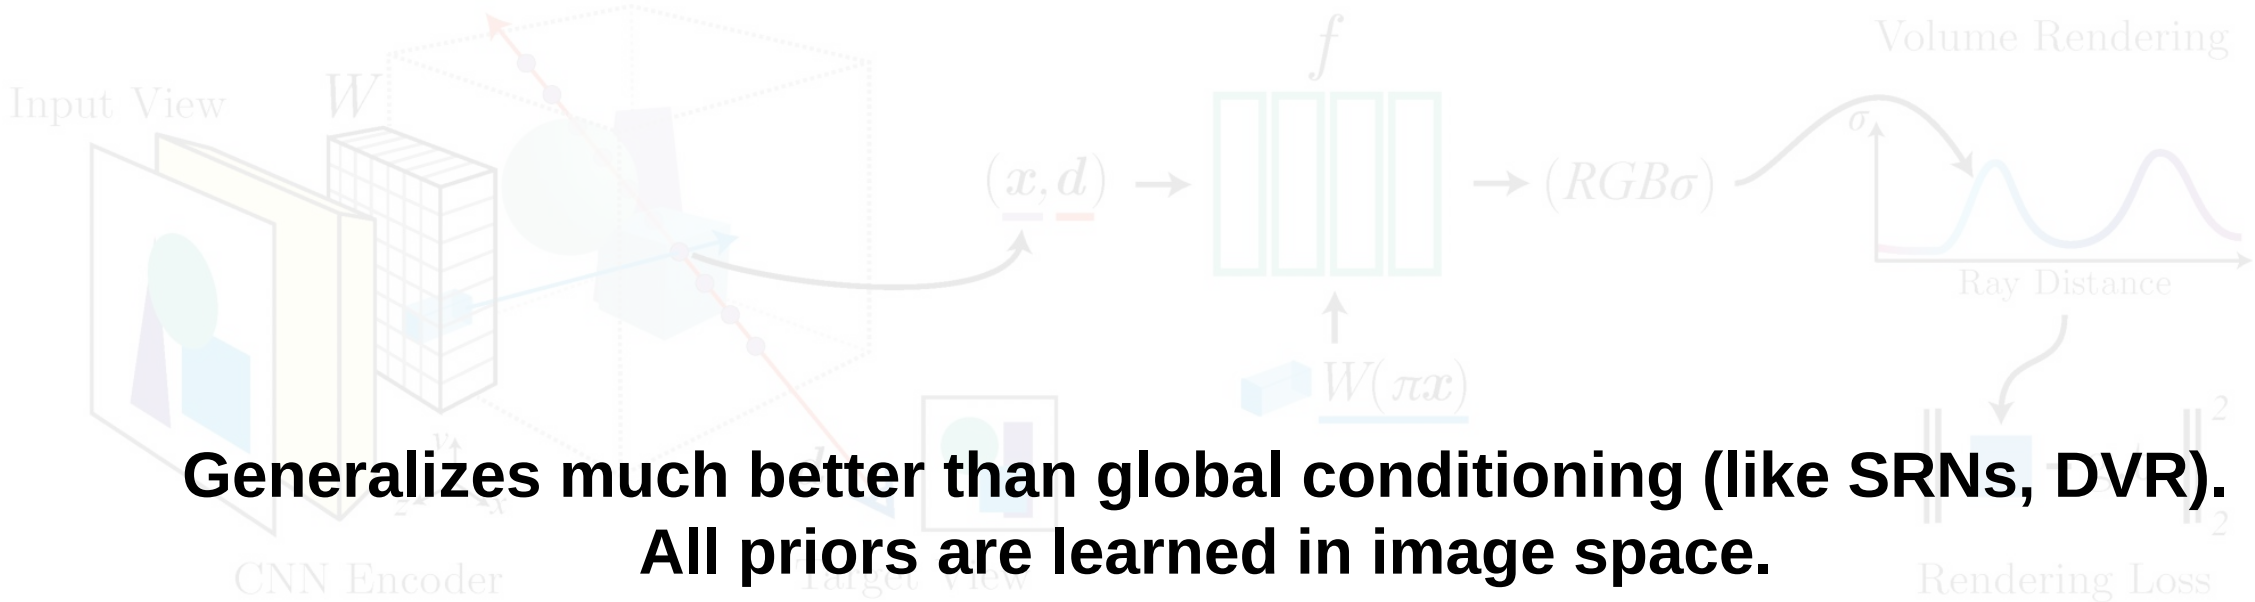

Generalizes much better than global conditioning (like SRNs, DVR).
All priors are learned in image space.

PiFU, Saito et al., ICCV 2019.
PixelNeRF, Yu et al., CVPR 2021
Grf: Learning a general radiance field…, Trevithick et al. ICCV 2021
Vision Transformer for NeRF-Based View Synthesis from a Single Input Image, Lin et al. 3DV 2022
MVSNeRF: Fast Generalizable Radiance Field Reconstruction from Multi-View Stereo, Chen et al. ICCV 2021
Neural Rays for Occlusion-aware Image-based Rendering, Liu et al. CVPR 2022
…

Object-centric representations

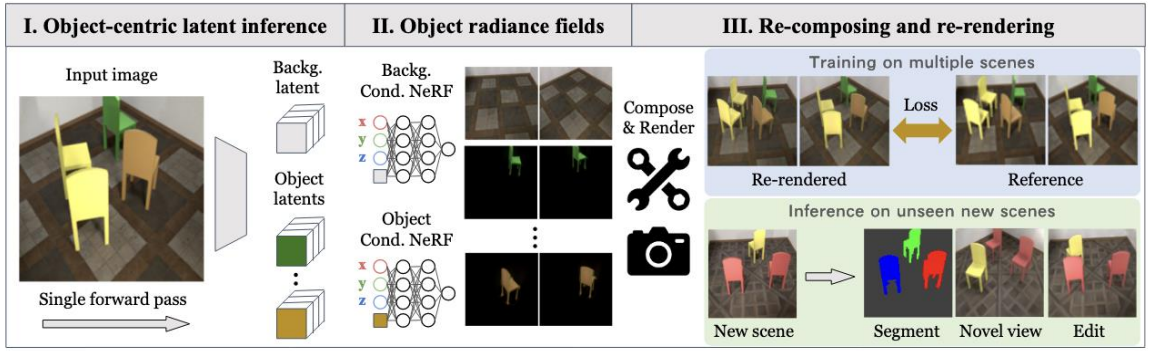

uORF, ICLR 2022

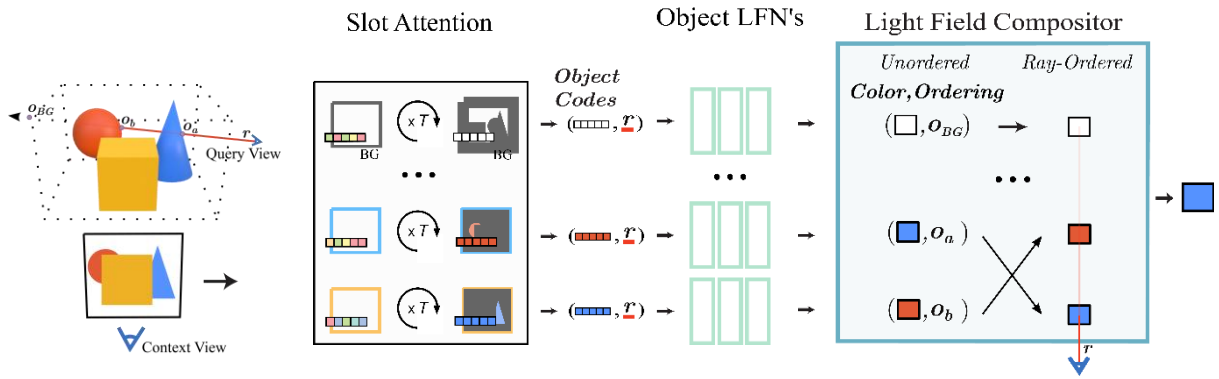

CoLF: Unsupervised Learning of Compositional Object Light Fields, arXiv 2022.

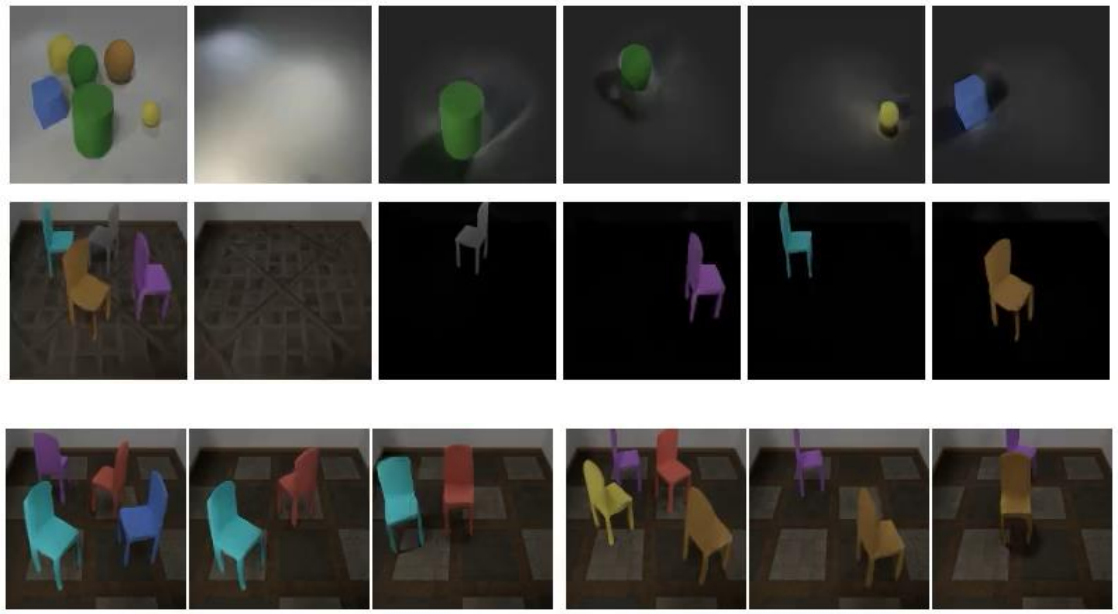

How to infer latent codes?

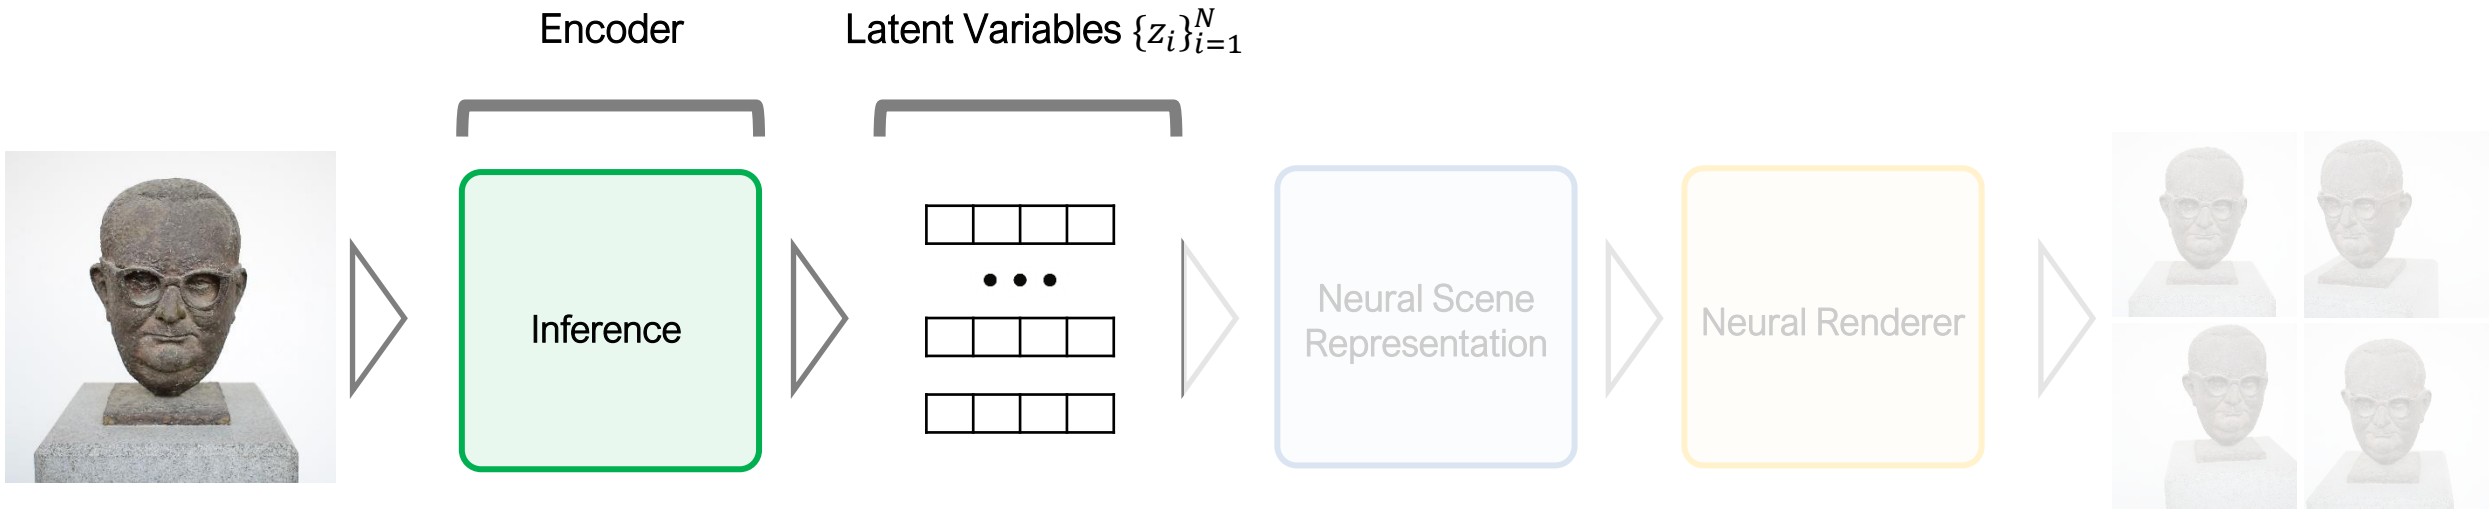

Encoding vs. Auto-Decoding

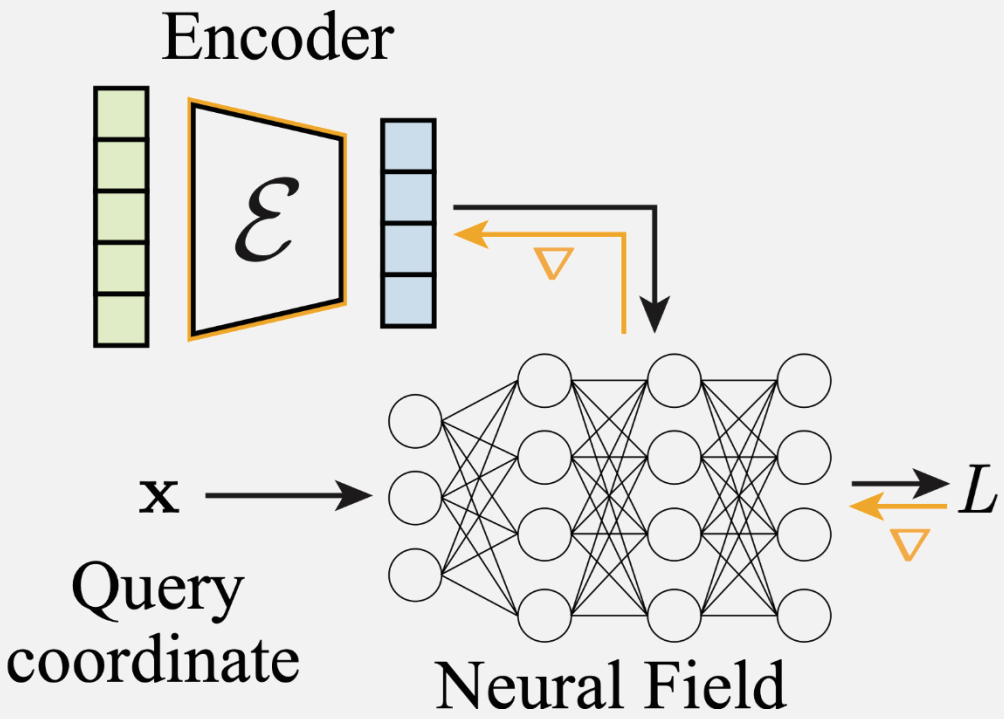

Encoding

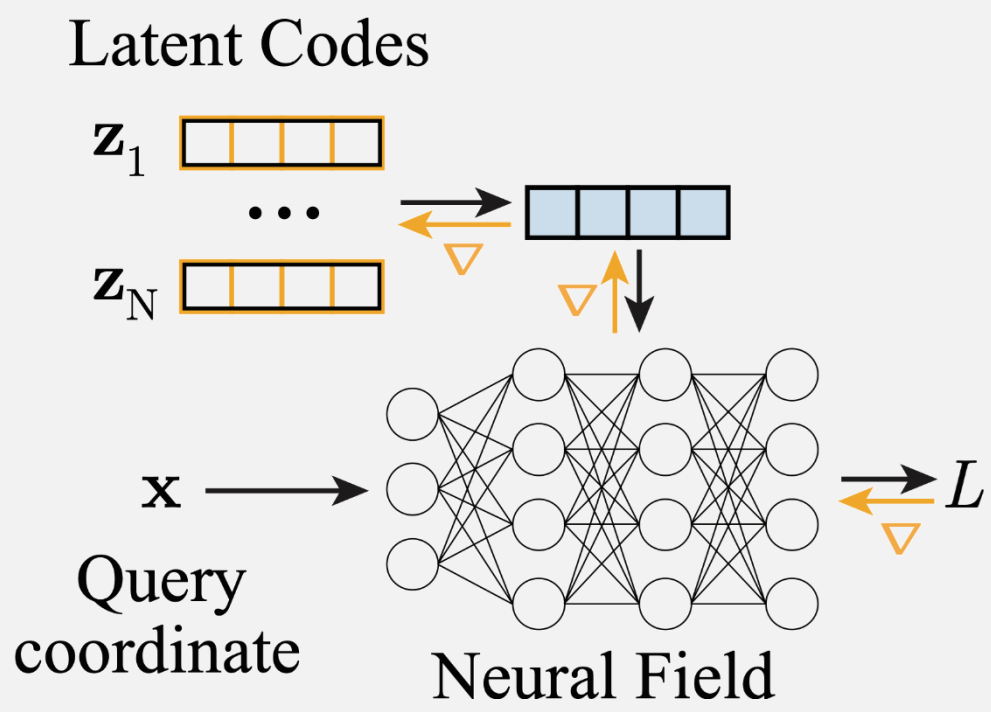

Auto-Decoding

Encoding

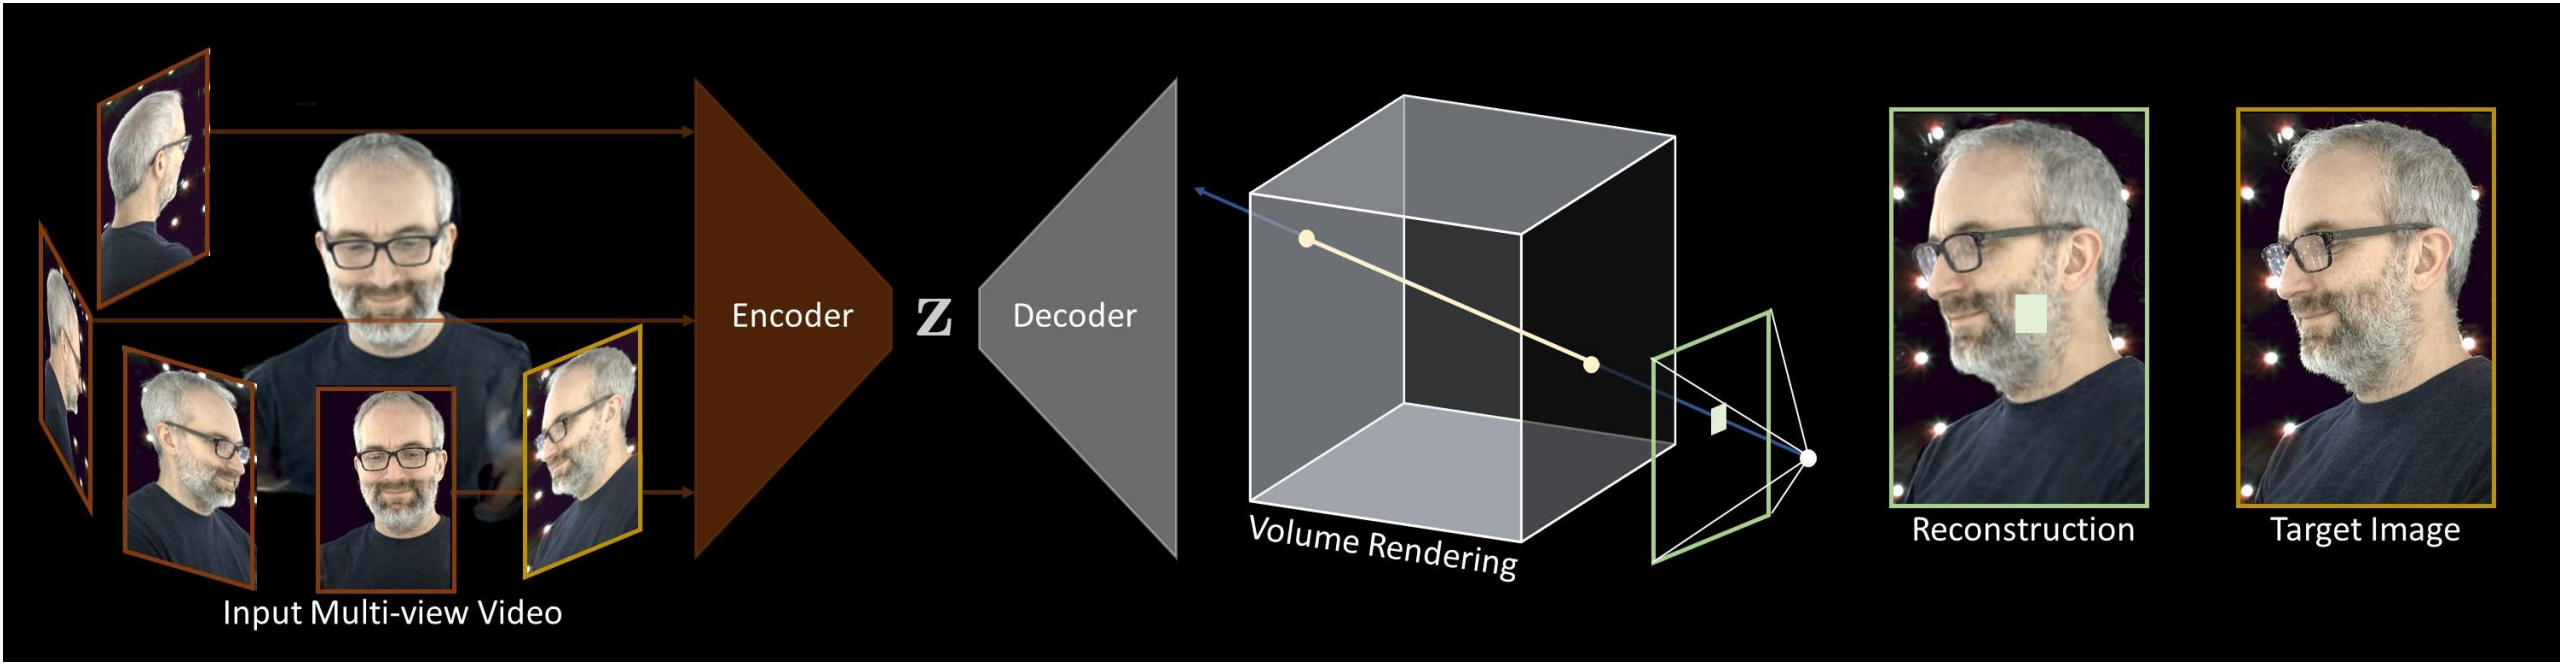

Note that, this is not to generalize to new objects, but generalize to different time frame of a dynamic scene

Neural Volumes, Lombodi et al., SIGGRAPH 2019

Encoding

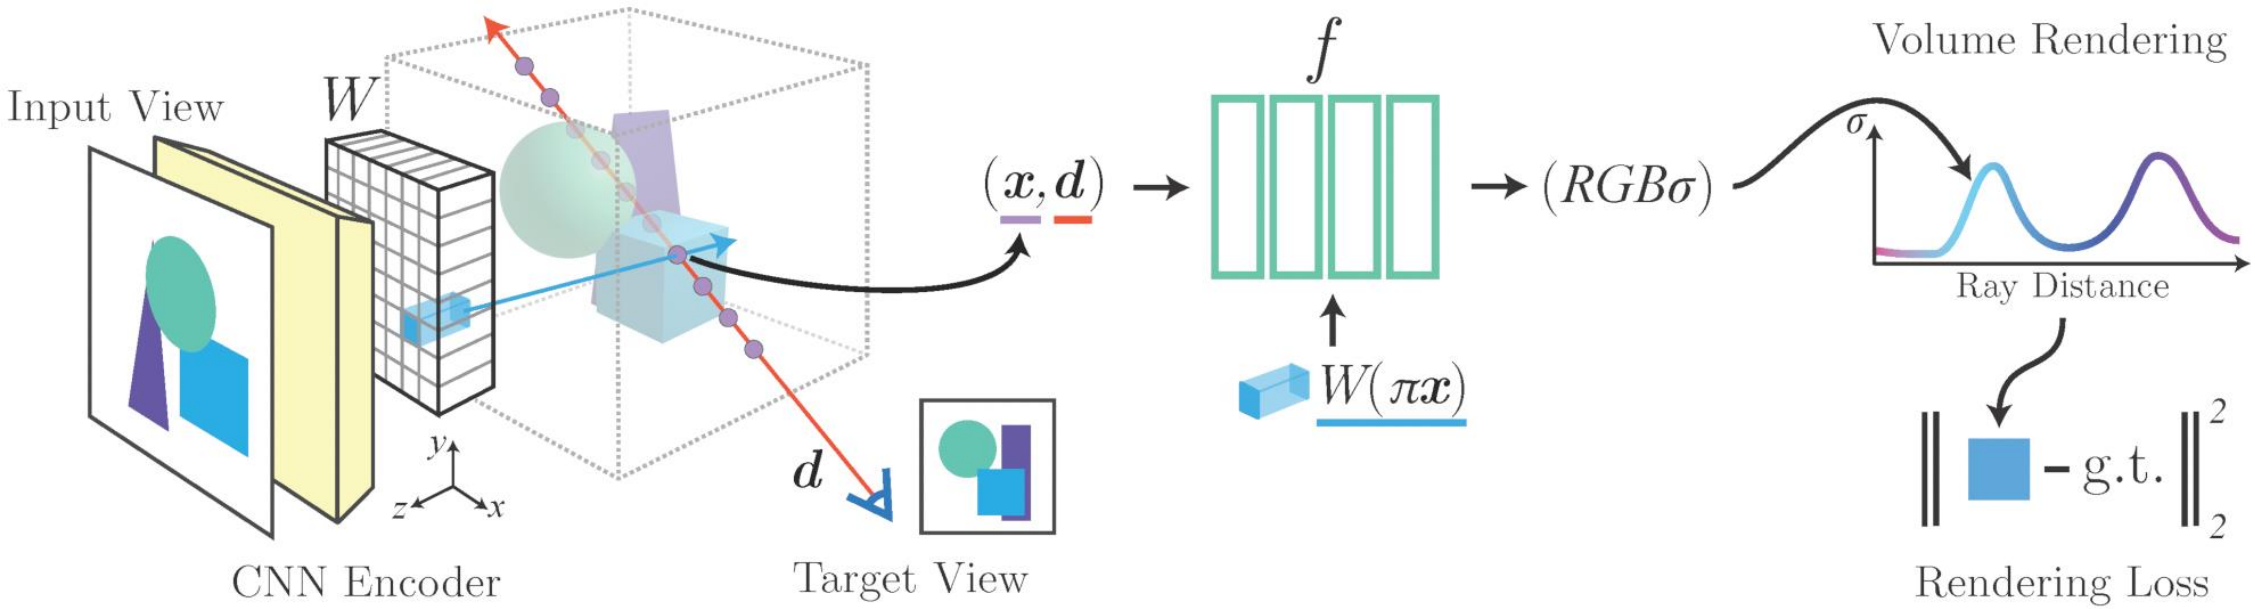

PixelNeRF, Yu et al., CVPR 2021

Auto-Decoding for inverse graphics

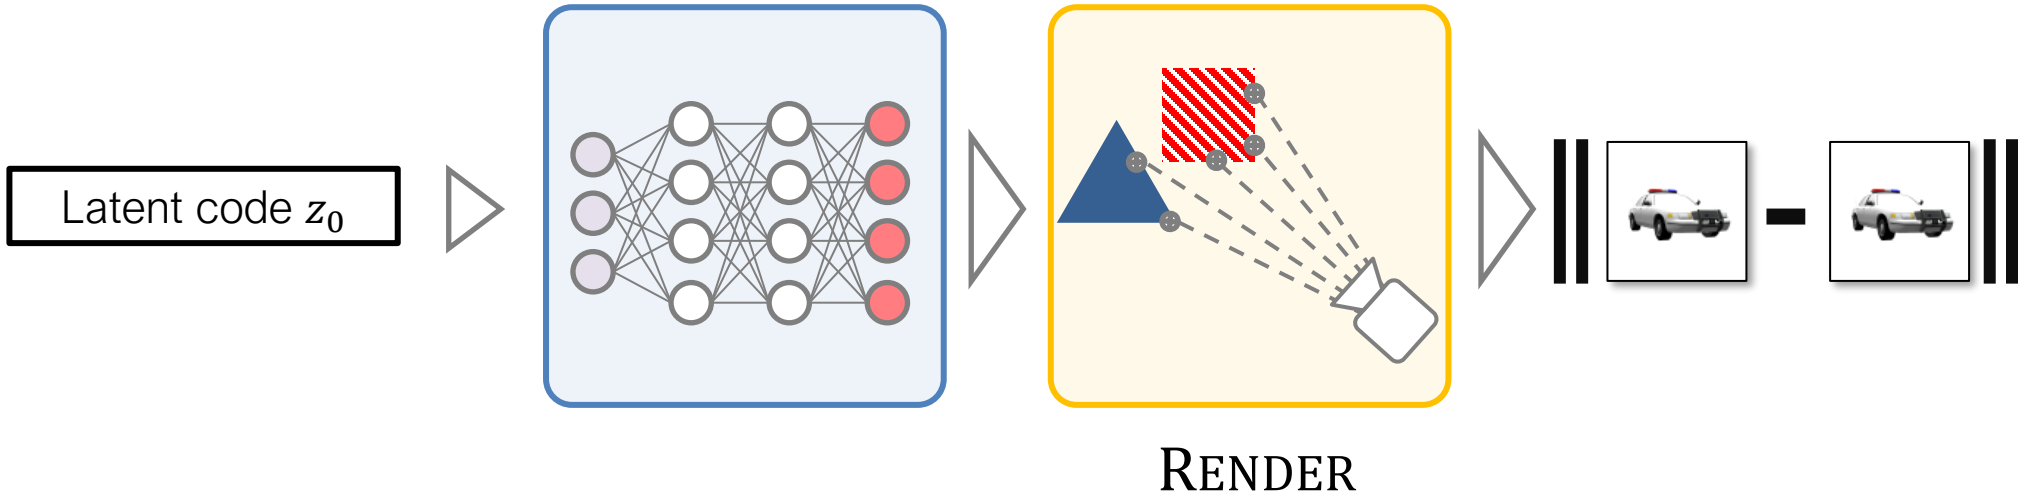

Sitzmann et al: Scene Representation Networks: Continuous 3D-Structure-Aware Neural Scene Representations, NeurIPS 2020.

Auto-Decoding for inverse graphics

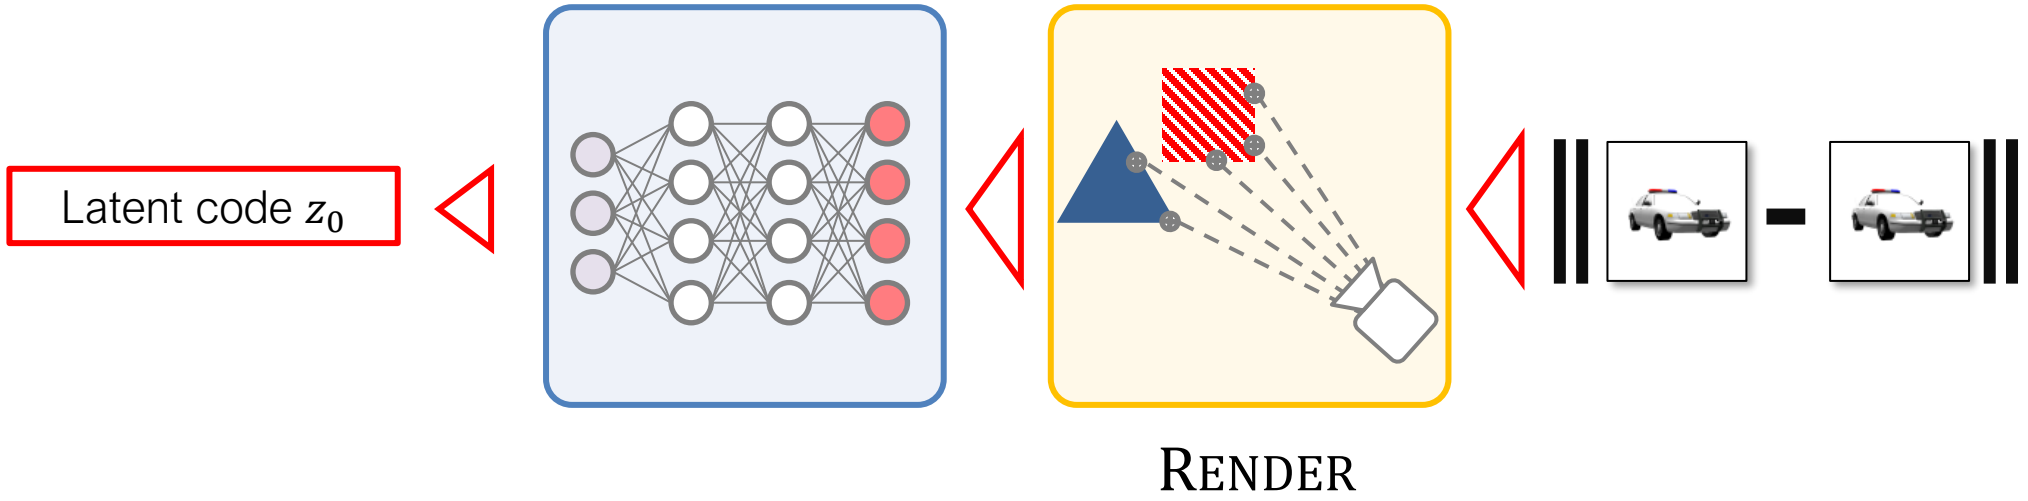

Sitzmann et al: Scene Representation Networks: Continuous 3D-Structure-Aware Neural Scene Representations, NeurIPS 2020.

Acknowledgments

- Vincent Sitzmann
- Neural Fields in Visual Computing and Beyond (Tutorial)
- Advances in Neural Rendering (Tutorial)
- awesome-NeRF: a curated list of awesome neural radiance fields papers
- MPII Summer Semester 2023: Computer Vision and Machine Learning for Computer Graphics
- Neural Volumetric Rendering for Computer Vision (Tutorial)