CoherentGS: Sparse Novel View Synthesis with 3D Gaussian Splatting

Avinash Paliwal[1,2], Wei Ye[2], Jinhui Xiong[2], Dmytro Kotovenko[3], Rakesh Ranjan[2], Vikas Chandra[2] and Nima Kalantari[1]

[1]Texas A&M University, [2]Meta Reality Labs, [3]LMU Munich

ECCV 2024

3D Gaussian Splatting (3DGS)

3D Gaussian Splatting (3DGS)

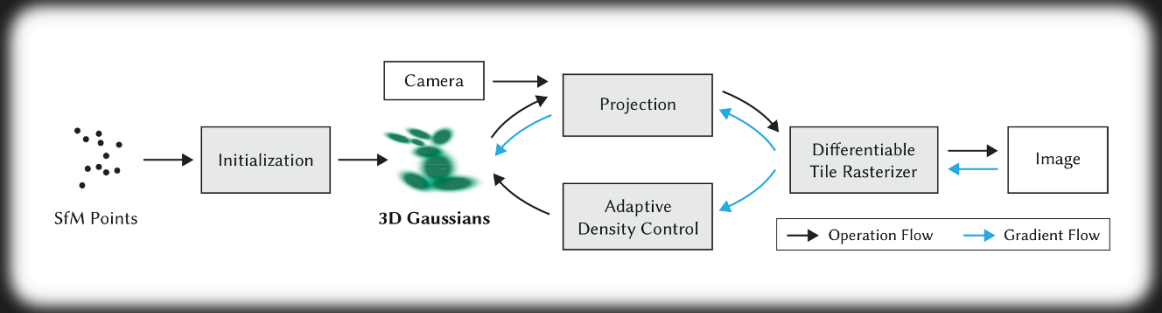

[Kerbl et al. 2023]

- High quality results with a large set of training images
- Real-time inference
- Overfits to sparse input

3DGS [Kerbl et al. 2023]

Goal

3D Gaussian Splatting (3DGS)

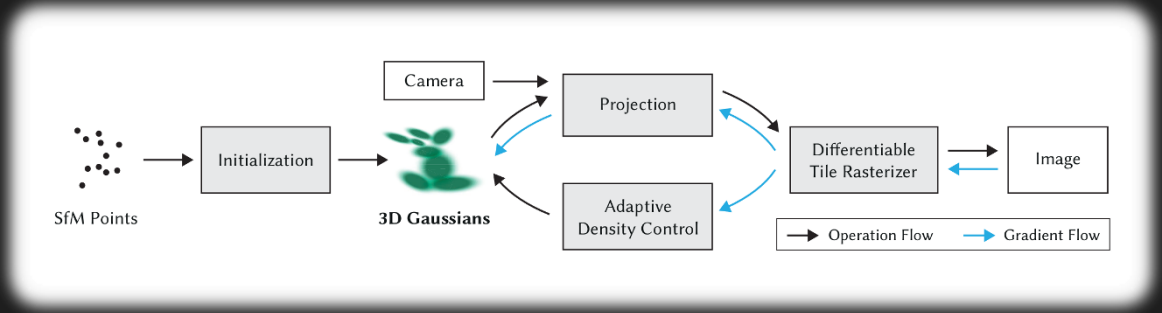

[Kerbl et al. 2023]

- Adapt 3DGS to extremely sparse input setting (2 – 4 images)

Ours

NeRF for Sparse Novel View Synthesis

Neural Radiance Field (NeRF)

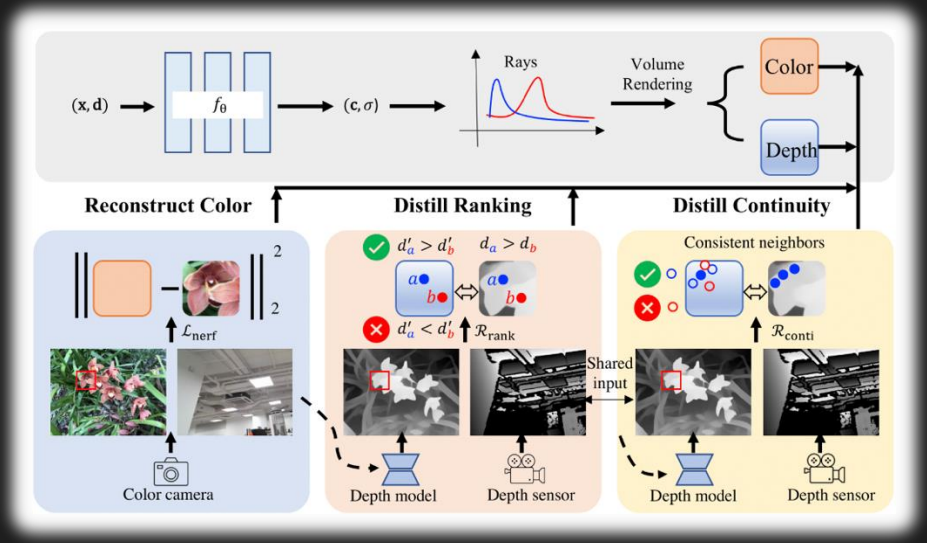

[Wang et al. 2023, Yang et al. 2023, Seo et al. 2023, Kangle et al. 2021, Niemeyer et al. 2021]

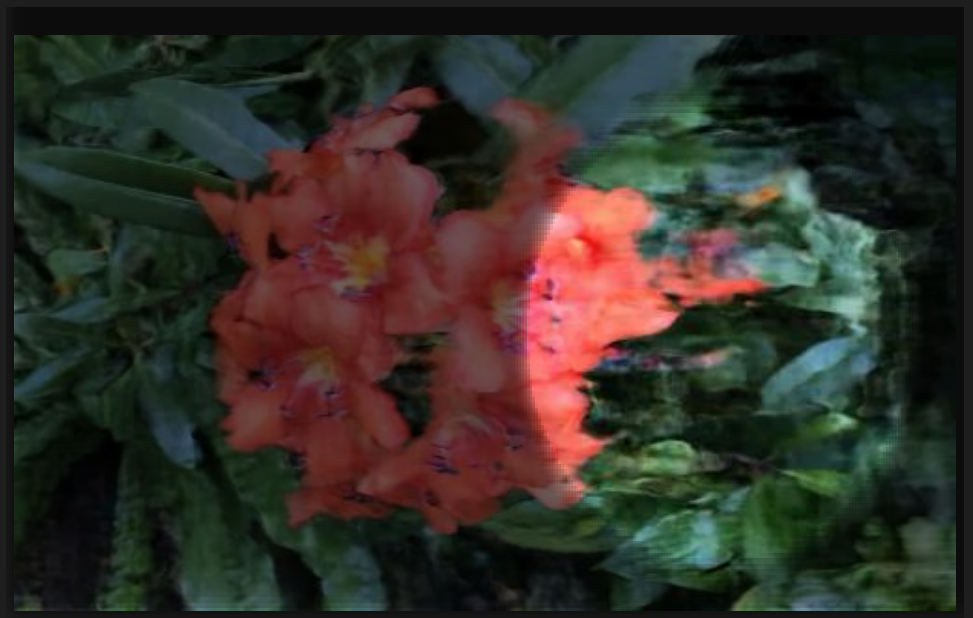

FreeNeRF [Yang et al. 2023]

- Additional regularizations to constrain the optimization
- Insufficient constraints for 2 to 4 inputs

NeRF for Sparse Novel View Synthesis

Neural Radiance Field (NeRF)

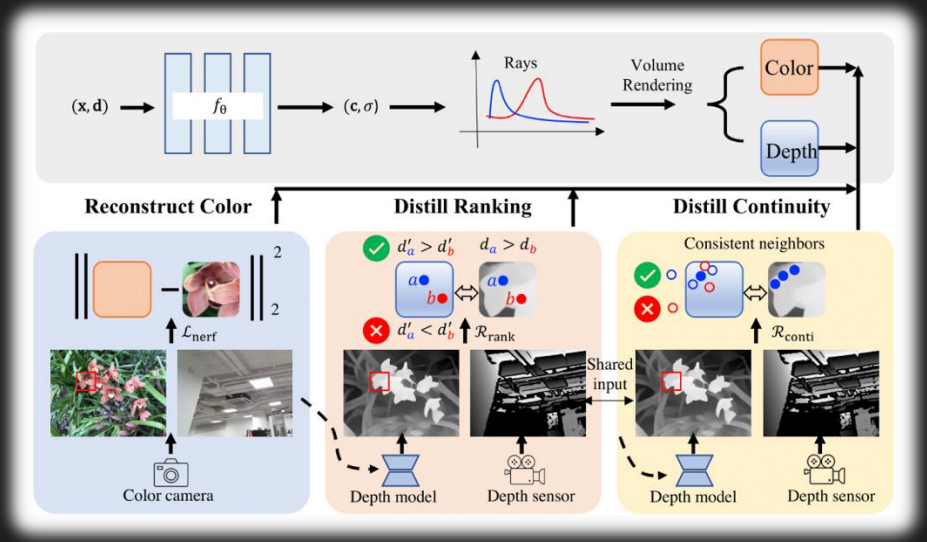

[Wang et al. 2023, Yang et al. 2023, Seo et al. 2023, Kangle et al. 2021, Niemeyer et al. 2021]

$$(x, y, z, \theta, \phi) \rightarrow F_\Theta \rightarrow (RGB\sigma)$$

NeRF for Sparse Novel View Synthesis

Neural Radiance Field (NeRF)

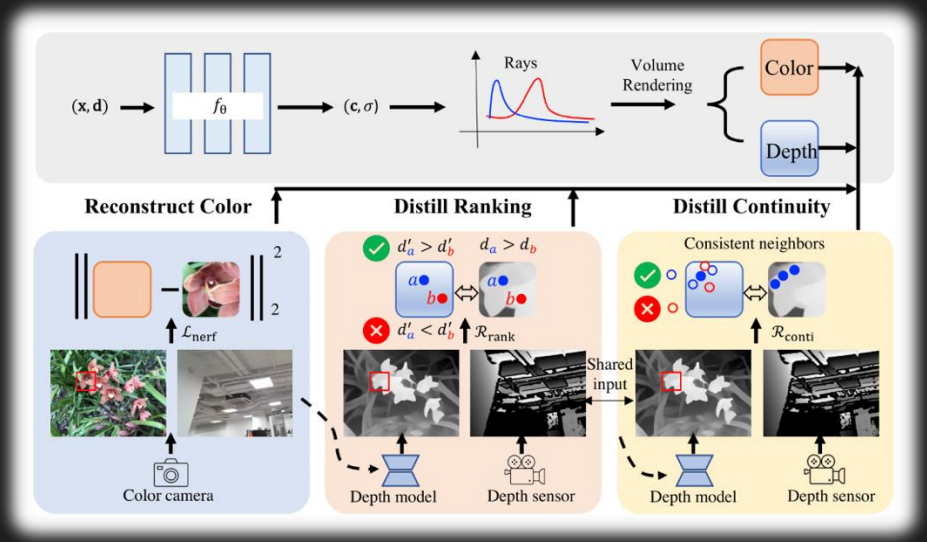

[Wang et al. 2023, Yang et al. 2023, Seo et al. 2023, Kangle et al. 2021, Niemeyer et al. 2021]

$$(x,y,z,\theta,\phi) \rightarrow F_\Theta \rightarrow (RGB\sigma)$$

3D Gaussian Splatting (3DGS)

3D Gaussian Splatting (3DGS)

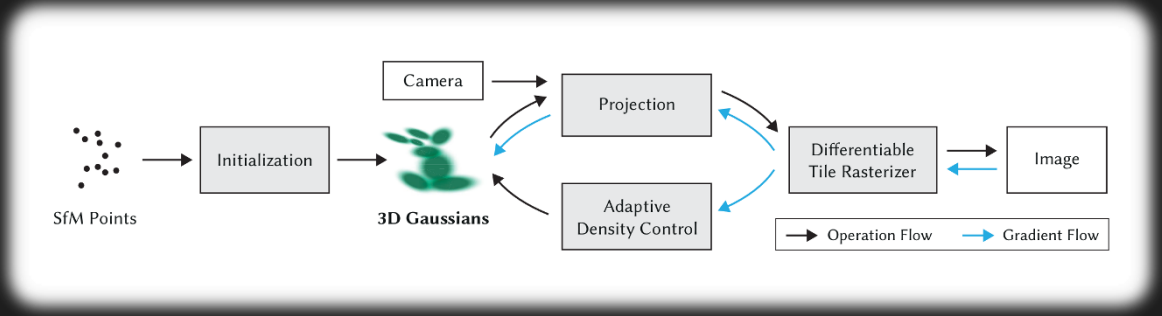

[Kerbl et al. 2023]

- Unstructured 3D representation
- Existing regularizations not as effective

3DGS [Kerbl et al. 2023]

Goal

3D Gaussian Splatting (3DGS)

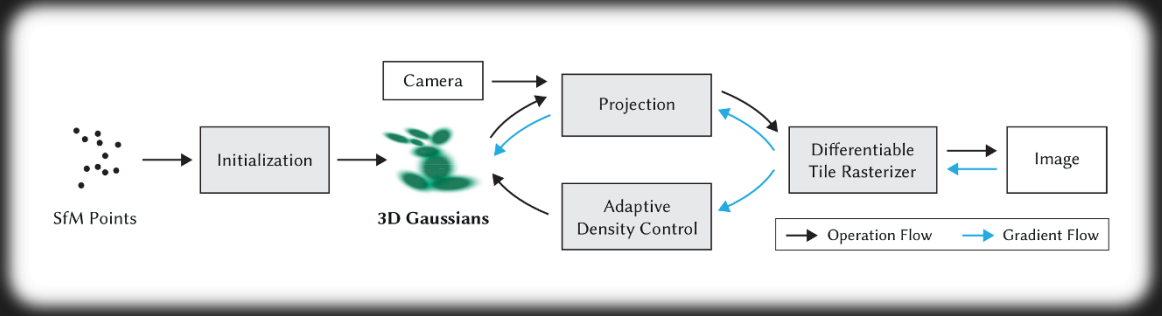

[Kerbl et al. 2023]

3DGS [Kerbl et al. 2023]

- Add coherency to the representation

Coherent Representation

- Difficult to enforce in the unstructured representation
- Add coherency in 2D domain

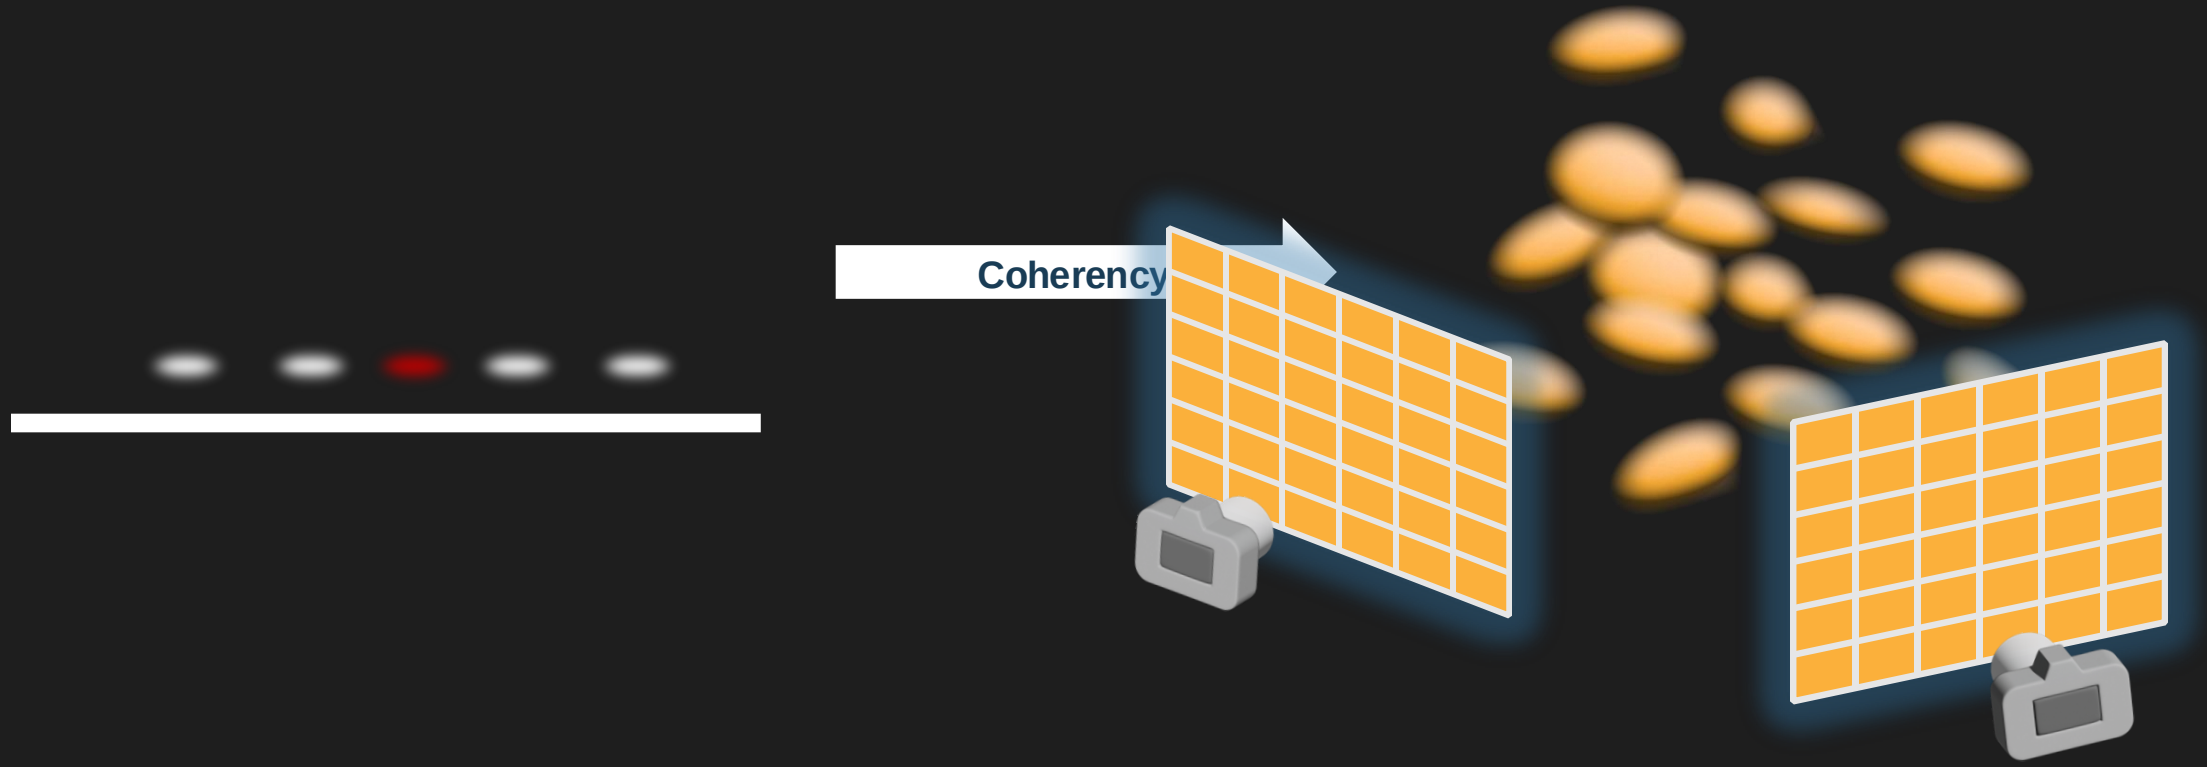

Coherent Representation

- Coherency in 2D domain
  - One gaussian per pixel
  - Gaussians move along camera ray

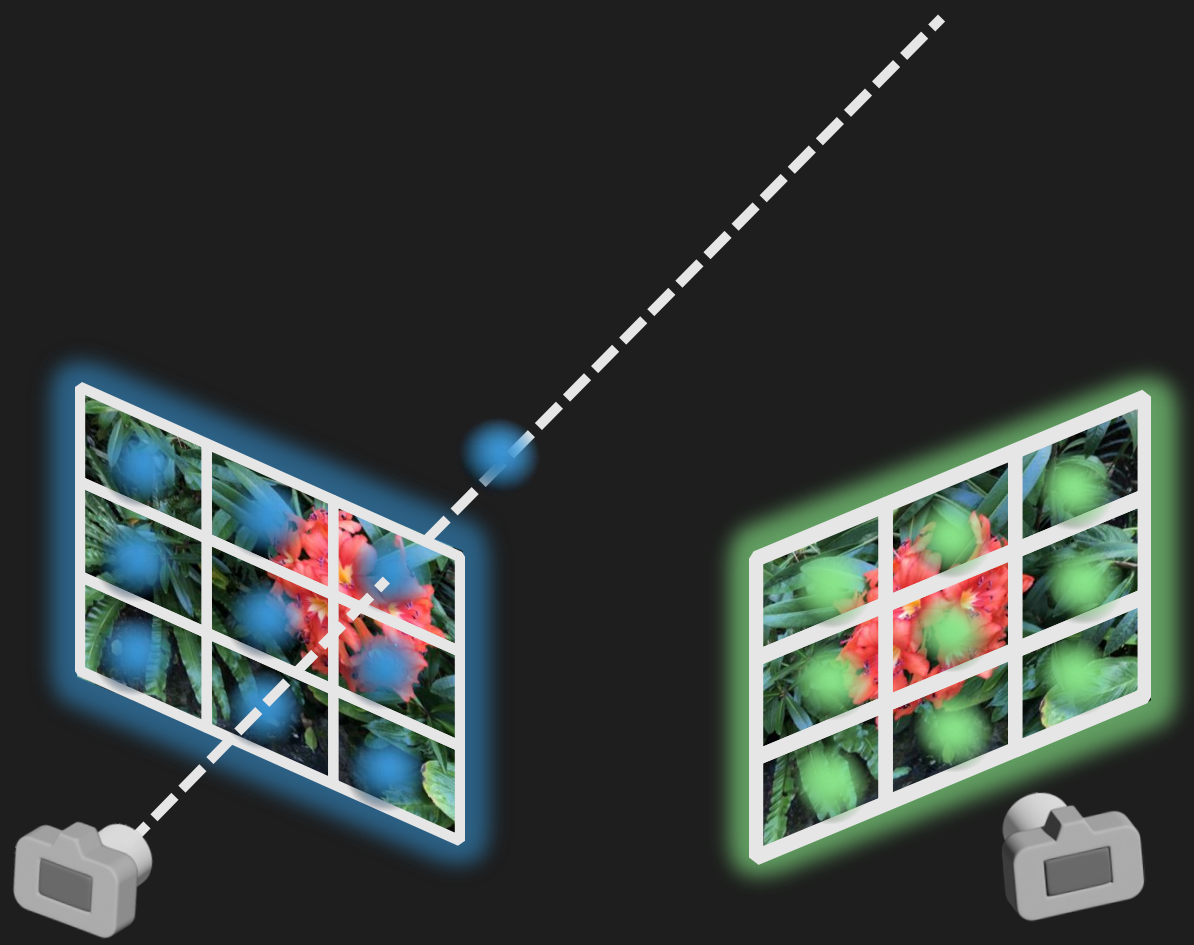

Coherent Representation

- Coherency in 2D domain
  - One gaussian per pixel
  - Gaussians move along camera ray
  - Gaussian position controlled by pixel depth
  - Enable regularizations in 2D

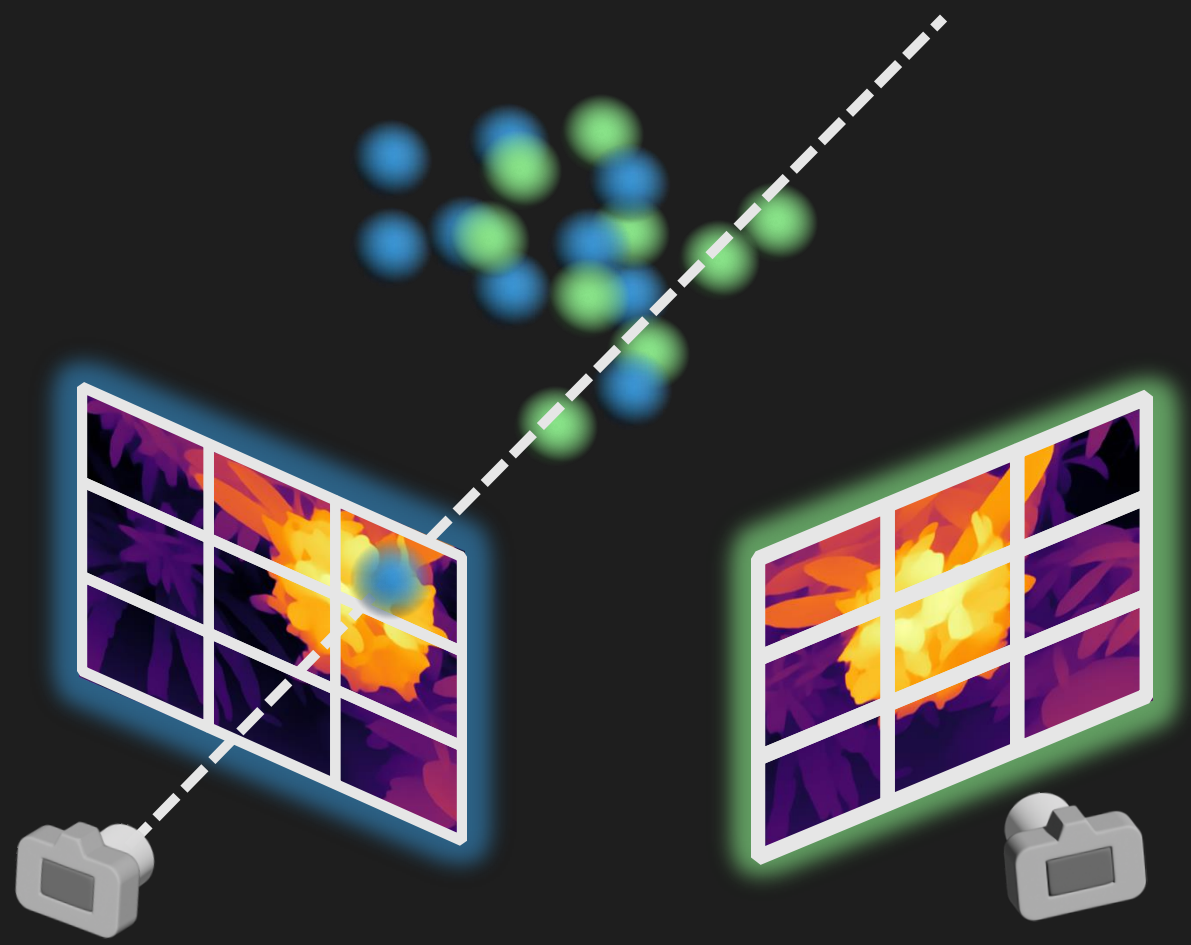

Constraints

- We add coherency to the 3D Gaussian representation
  - Single-view constraints
  - Multi-view constraints

Single-view Constraints

- Coherency in 2D domain

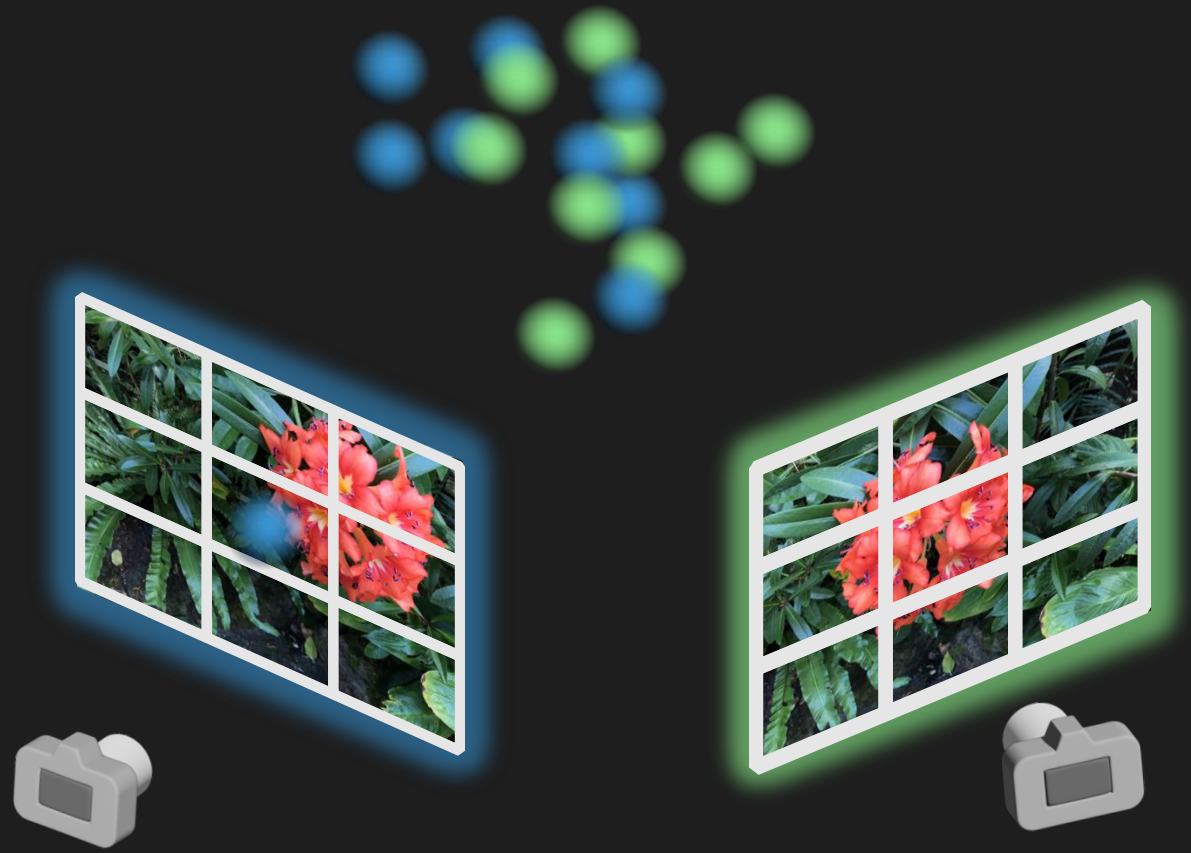

Single-view Constraints

- Coherency in 2D domain
  - Implicit decoder predicts smooth 2D depth
  - Gaussians move together as a group

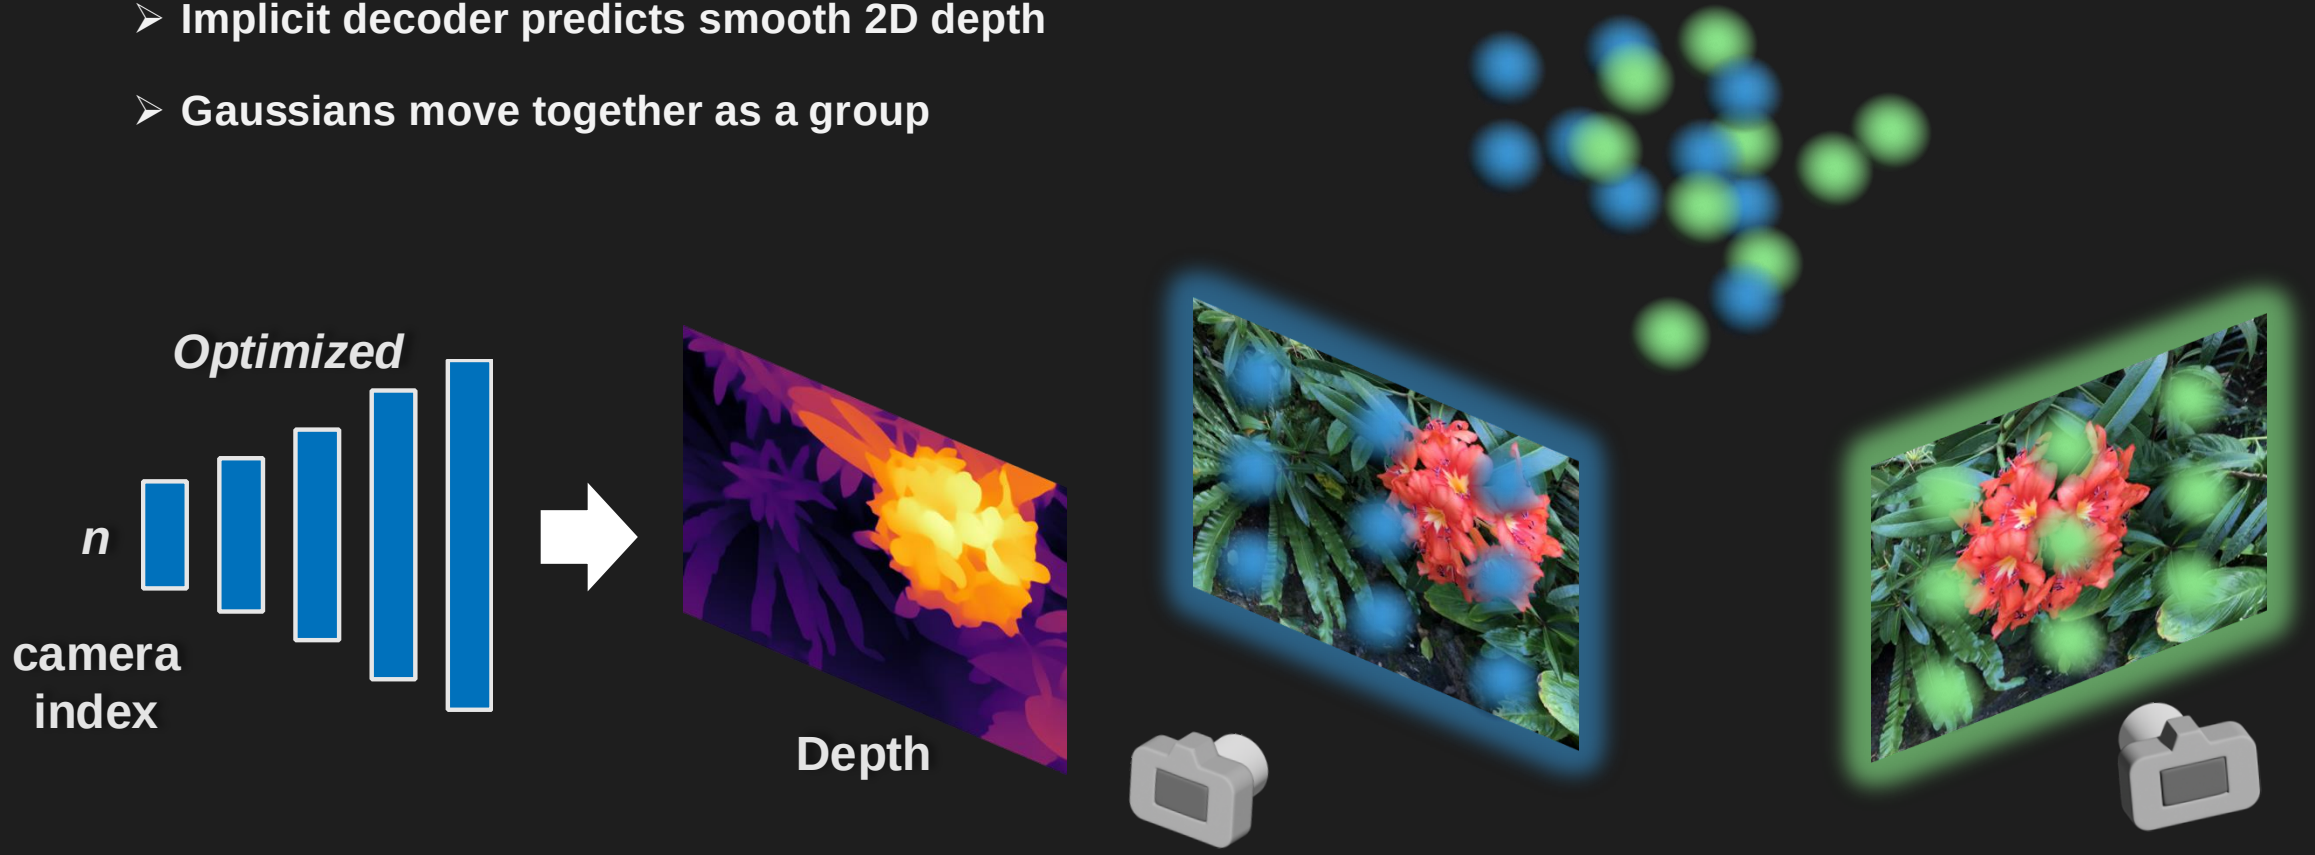

Multi-view Constraints

- Coherency in 2D domain
  - Rasterize all the gaussians to obtain scene depth
  - Apply total variation regularization

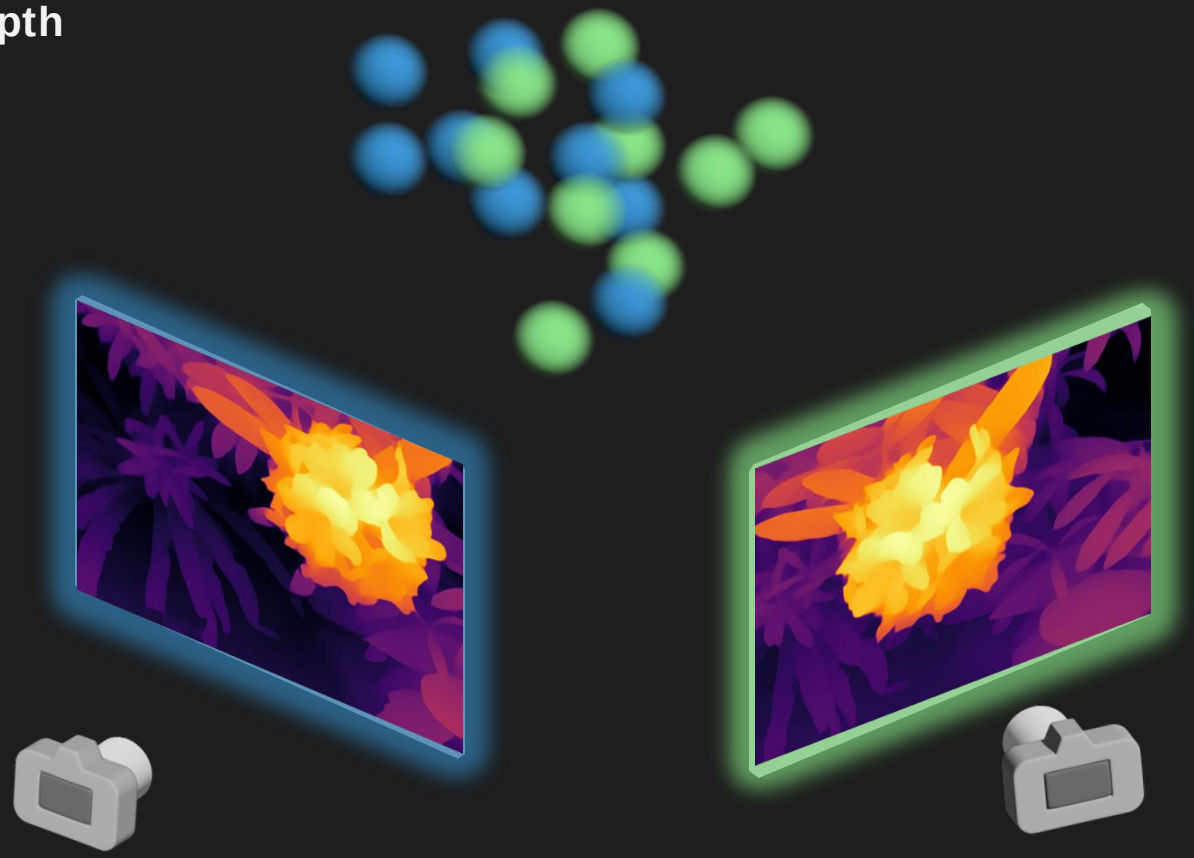

Optimization

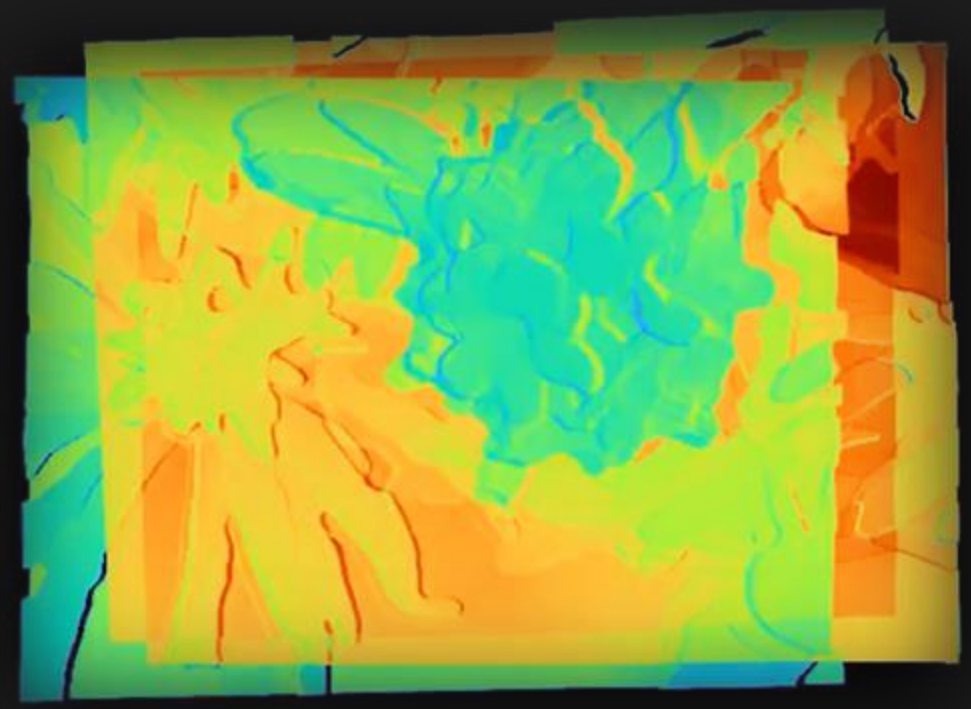

- Gaussians from different views are quickly aligned due to coherent movement
- Obtain high quality geometry and texture without noisy gaussians

Inpainting Results – 3 inputs

Ours w/o inpainting

Ours w/ inpainting

Inpainting Results – 3 inputs

FreeNeRF [Yang et al. 2023]

SparseNeRF [Wang et al. 2023]

Ours w/o inpainting

Ours w/ inpainting

Results – 2 inputs

FlipNeRF [Seo et al. 2023]

SparseNeRF [Wang et al. 2023]

FreeNeRF [Yang et al. 2023]

Ours

Results – 3 inputs

FlipNeRF [Seo et al. 2023]

SparseNeRF [Wang et al. 2023]

FreeNeRF [Yang et al. 2023]

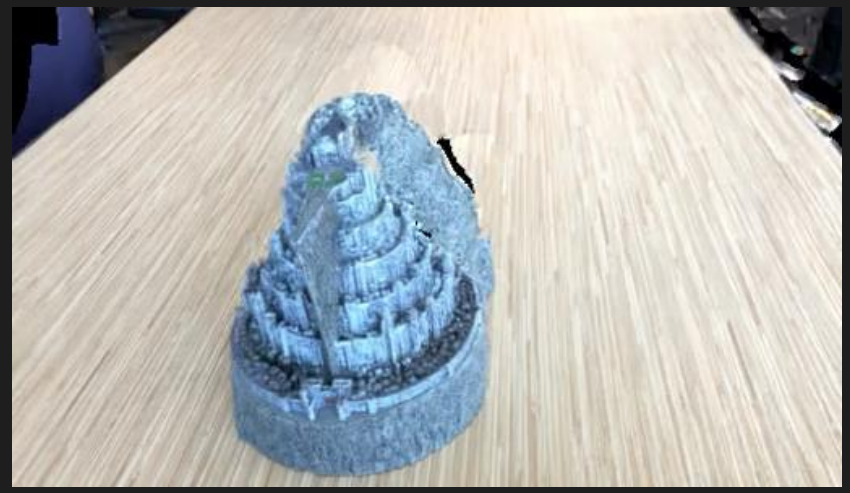

Ours

Thank You

Poster #319

Session: Thu 3 Oct 10:30 a.m. - 12:30 a.m.