nerf2nerf: Pairwise Registration of Neural Radiance Fields

Lily Goli[1,2], Daniel Rebain[4], Sara Sabour[1,2,6], Animesh Garg[1,2,5], Andrea Tagliasacchi[1,3,6]

[1]University of Toronto, [2]Vector Institute, [3]SFU, [4]UBC, [5]NVIDIA, [6]Google Research

Registration Applications

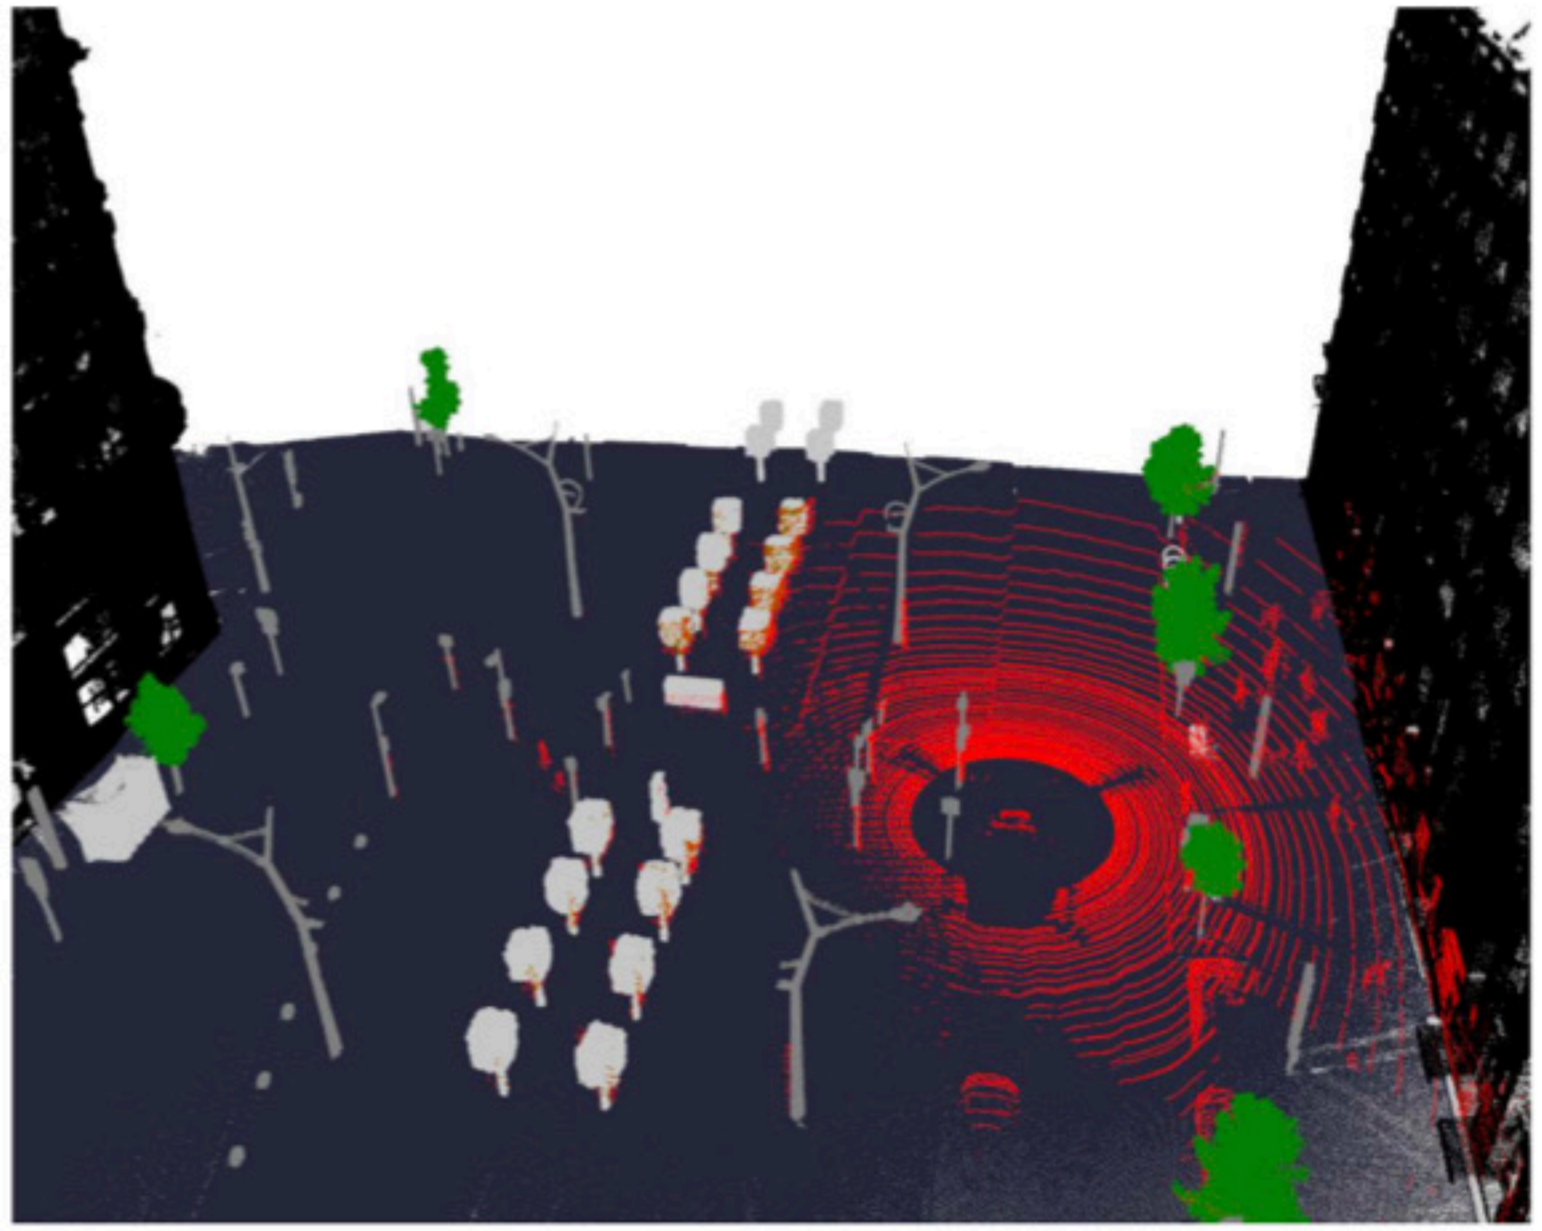

Nagy et al. ECCV'18

"localization"

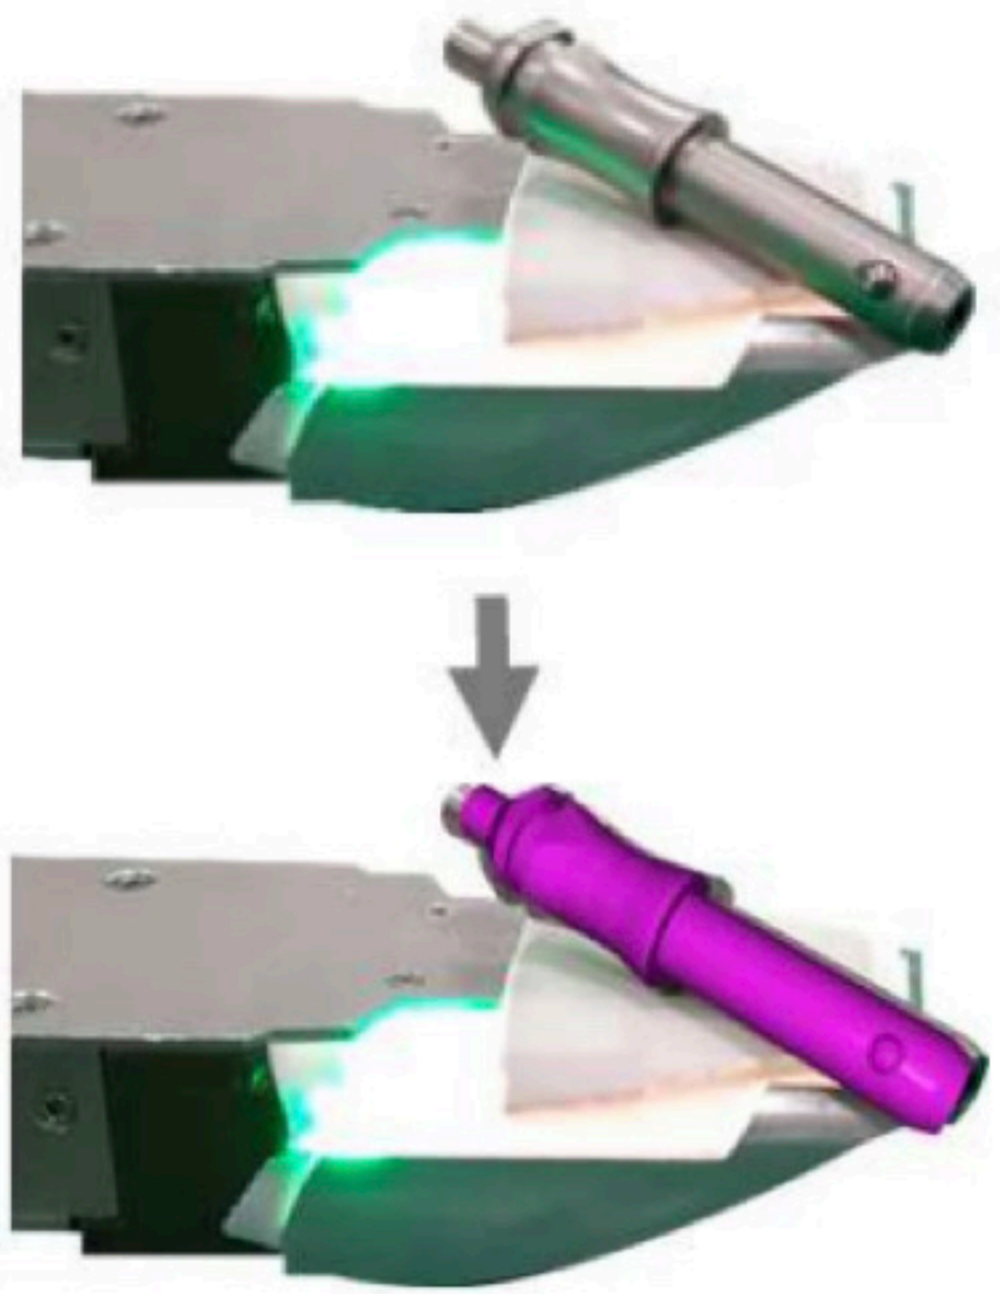

Bauzá et al. CoRL'20

"Pose Estimation"

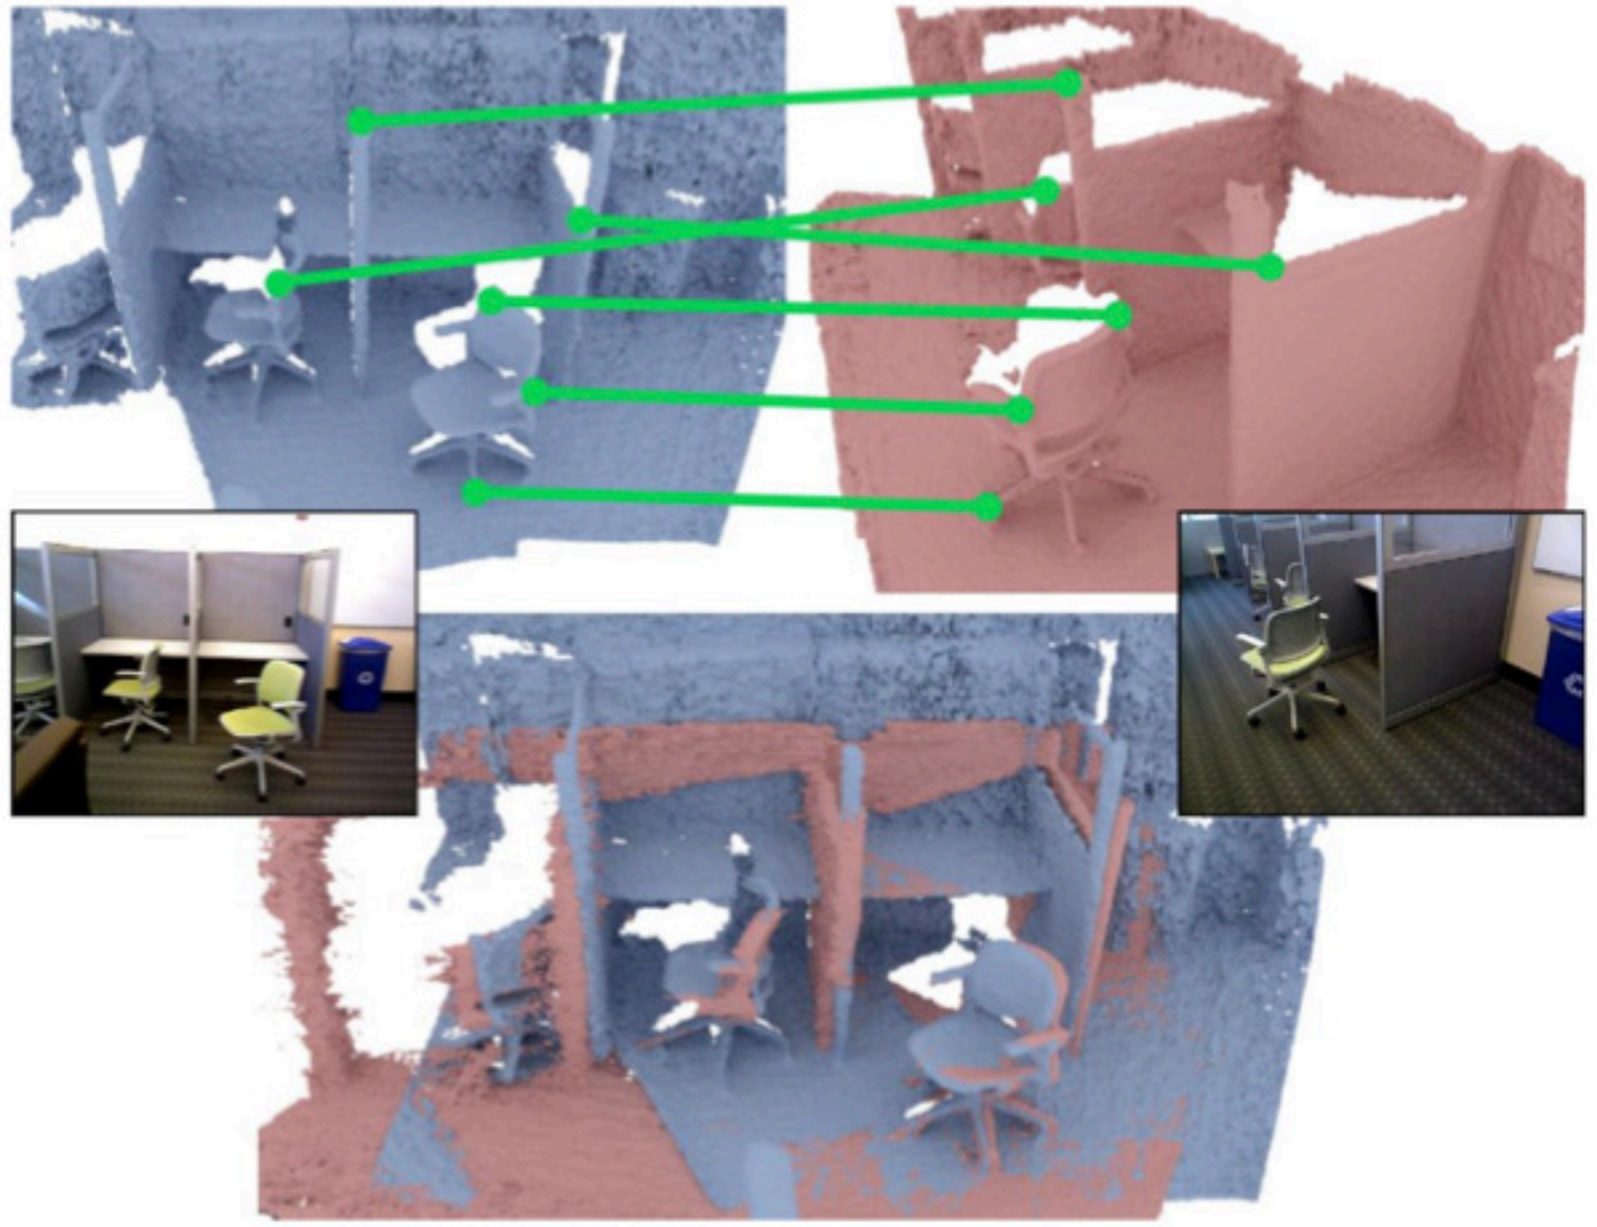

Zeng et al. CVPR'17

"Large Scene Reconstruction"

Lily Goli, Daniel Rebain, Sara Sabour, Animesh Garg, Andrea Tagliasacchi "nerf2nerf:Pairwise Registration of Neural Radiance Fields" ICRA 2023

Classic 3D Registration

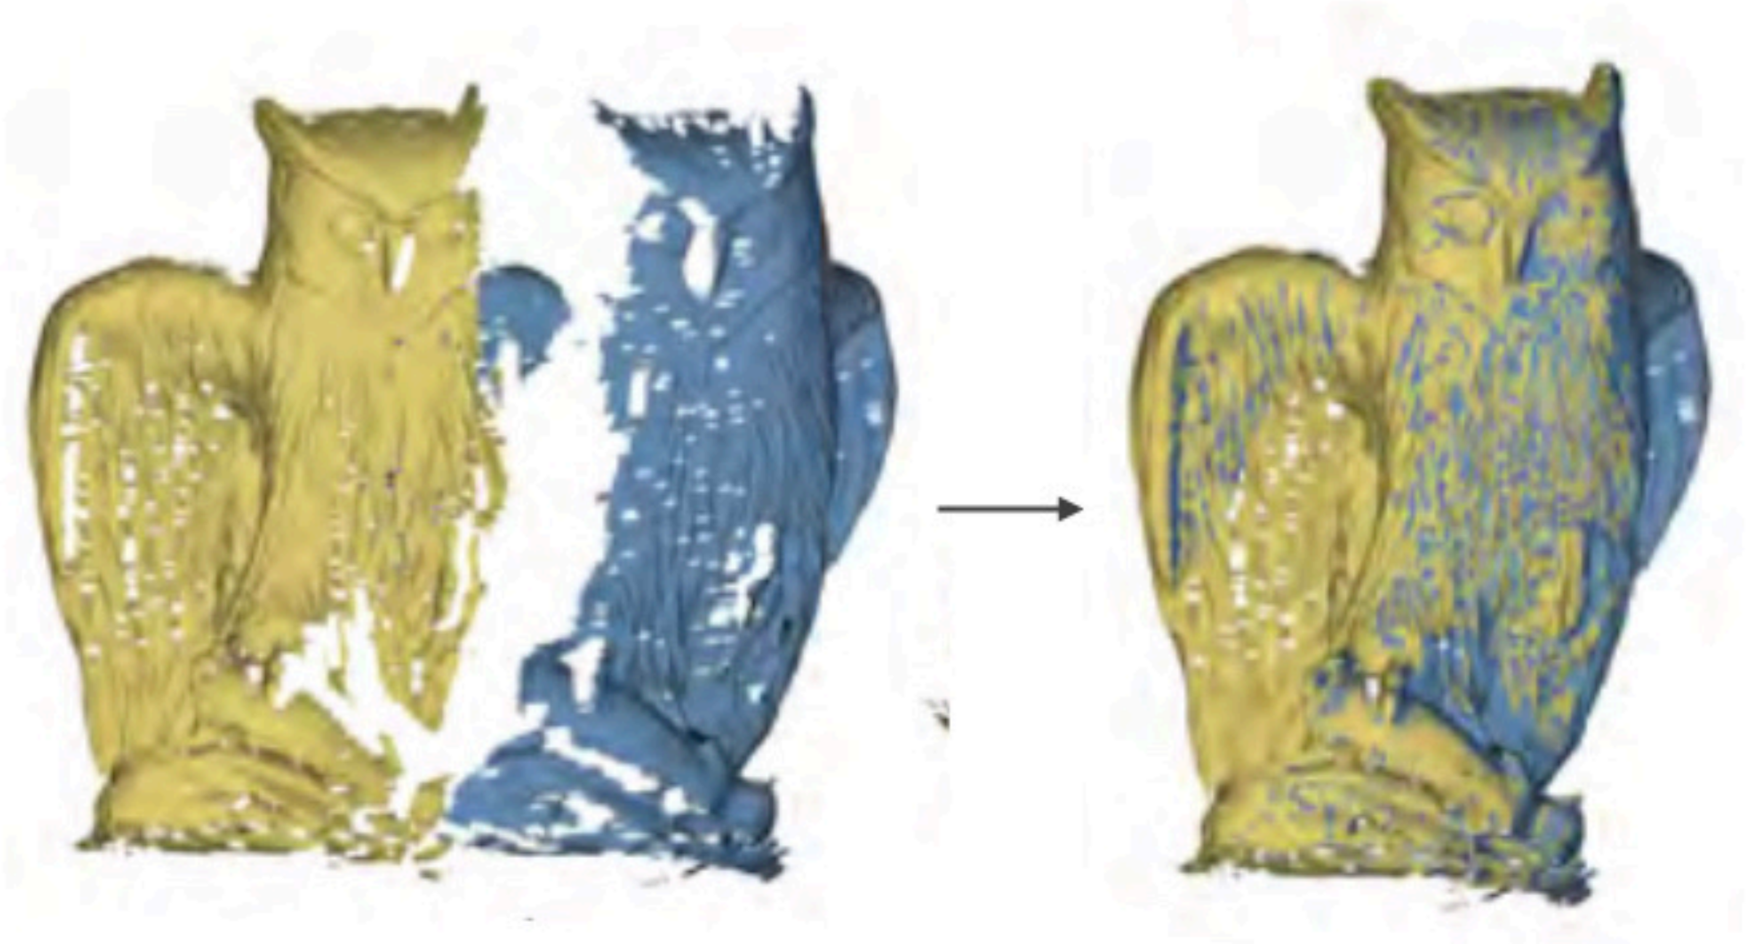

Bouaziz et al. SGP'13
Sparse Iterative Closest Point
(Sparse ICP)

Zhou et al. ECCV'16
Fast Global Registration (FGR)

Lily Goli, Daniel Rebain, Sara Sabour, Animesh Garg, Andrea Tagliasacchi "nerf2nerf:Pairwise Registration of Neural Radiance Fields" ICRA 2023

Neural Radiance Field (NeRF)

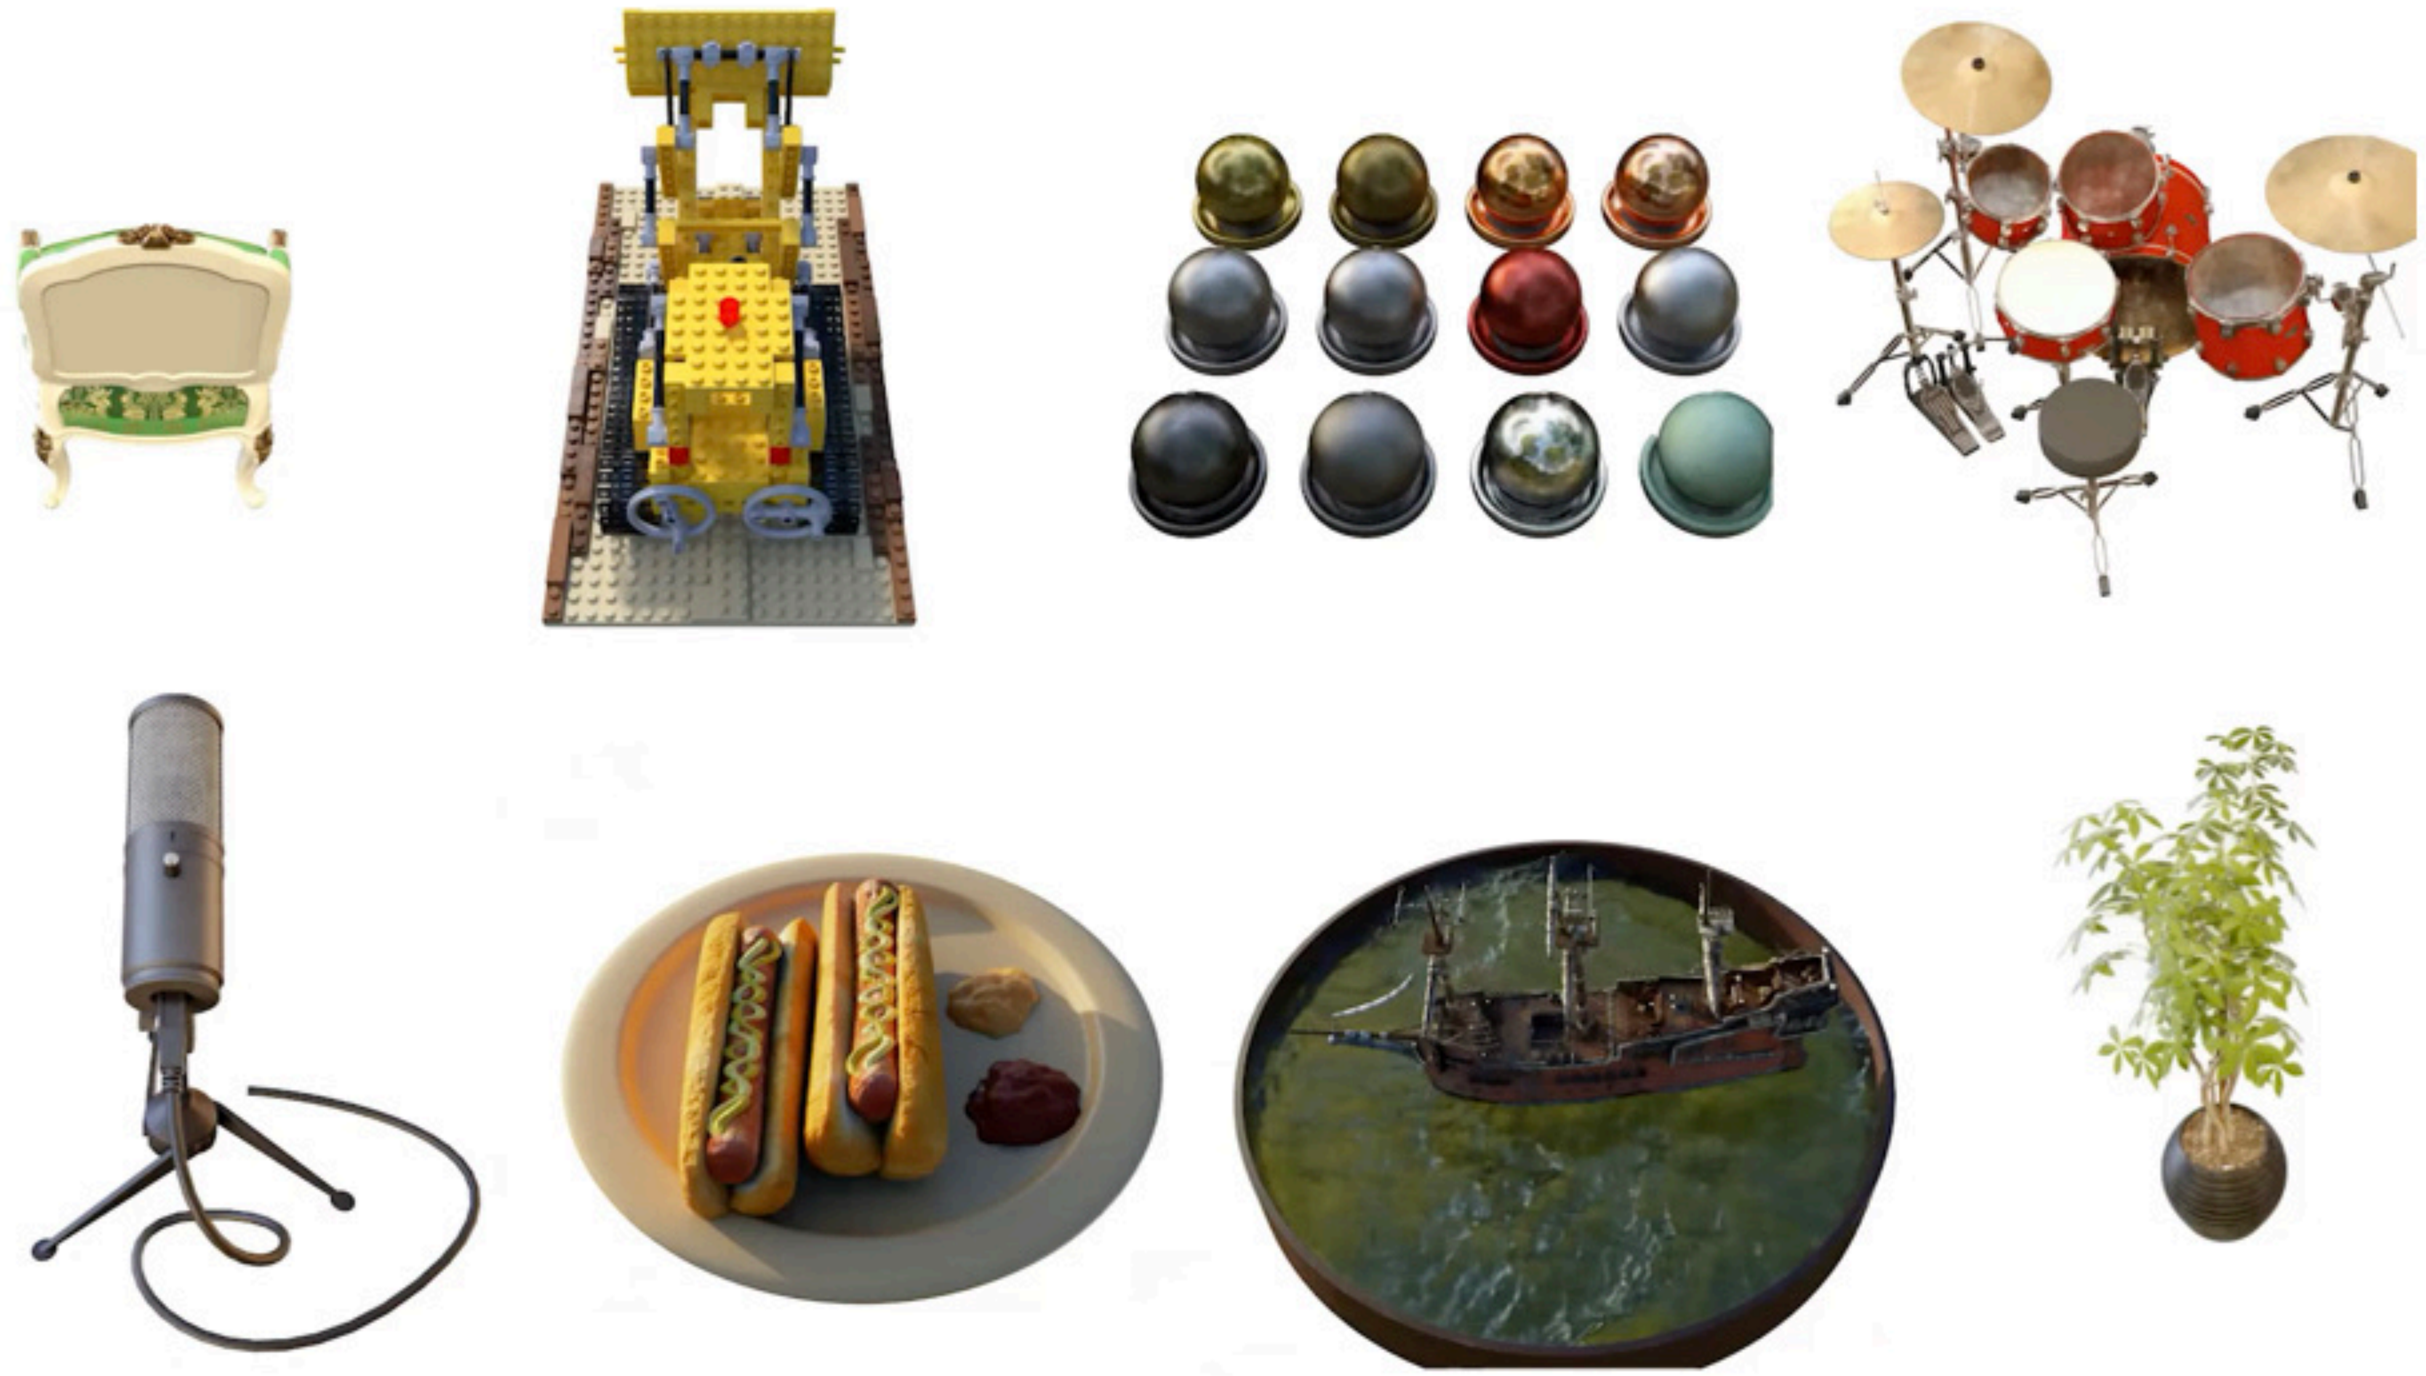

Mildenhall, Srinivasan, Tanick et al. ECCV'20
Neural Radiance Field (NeRF)

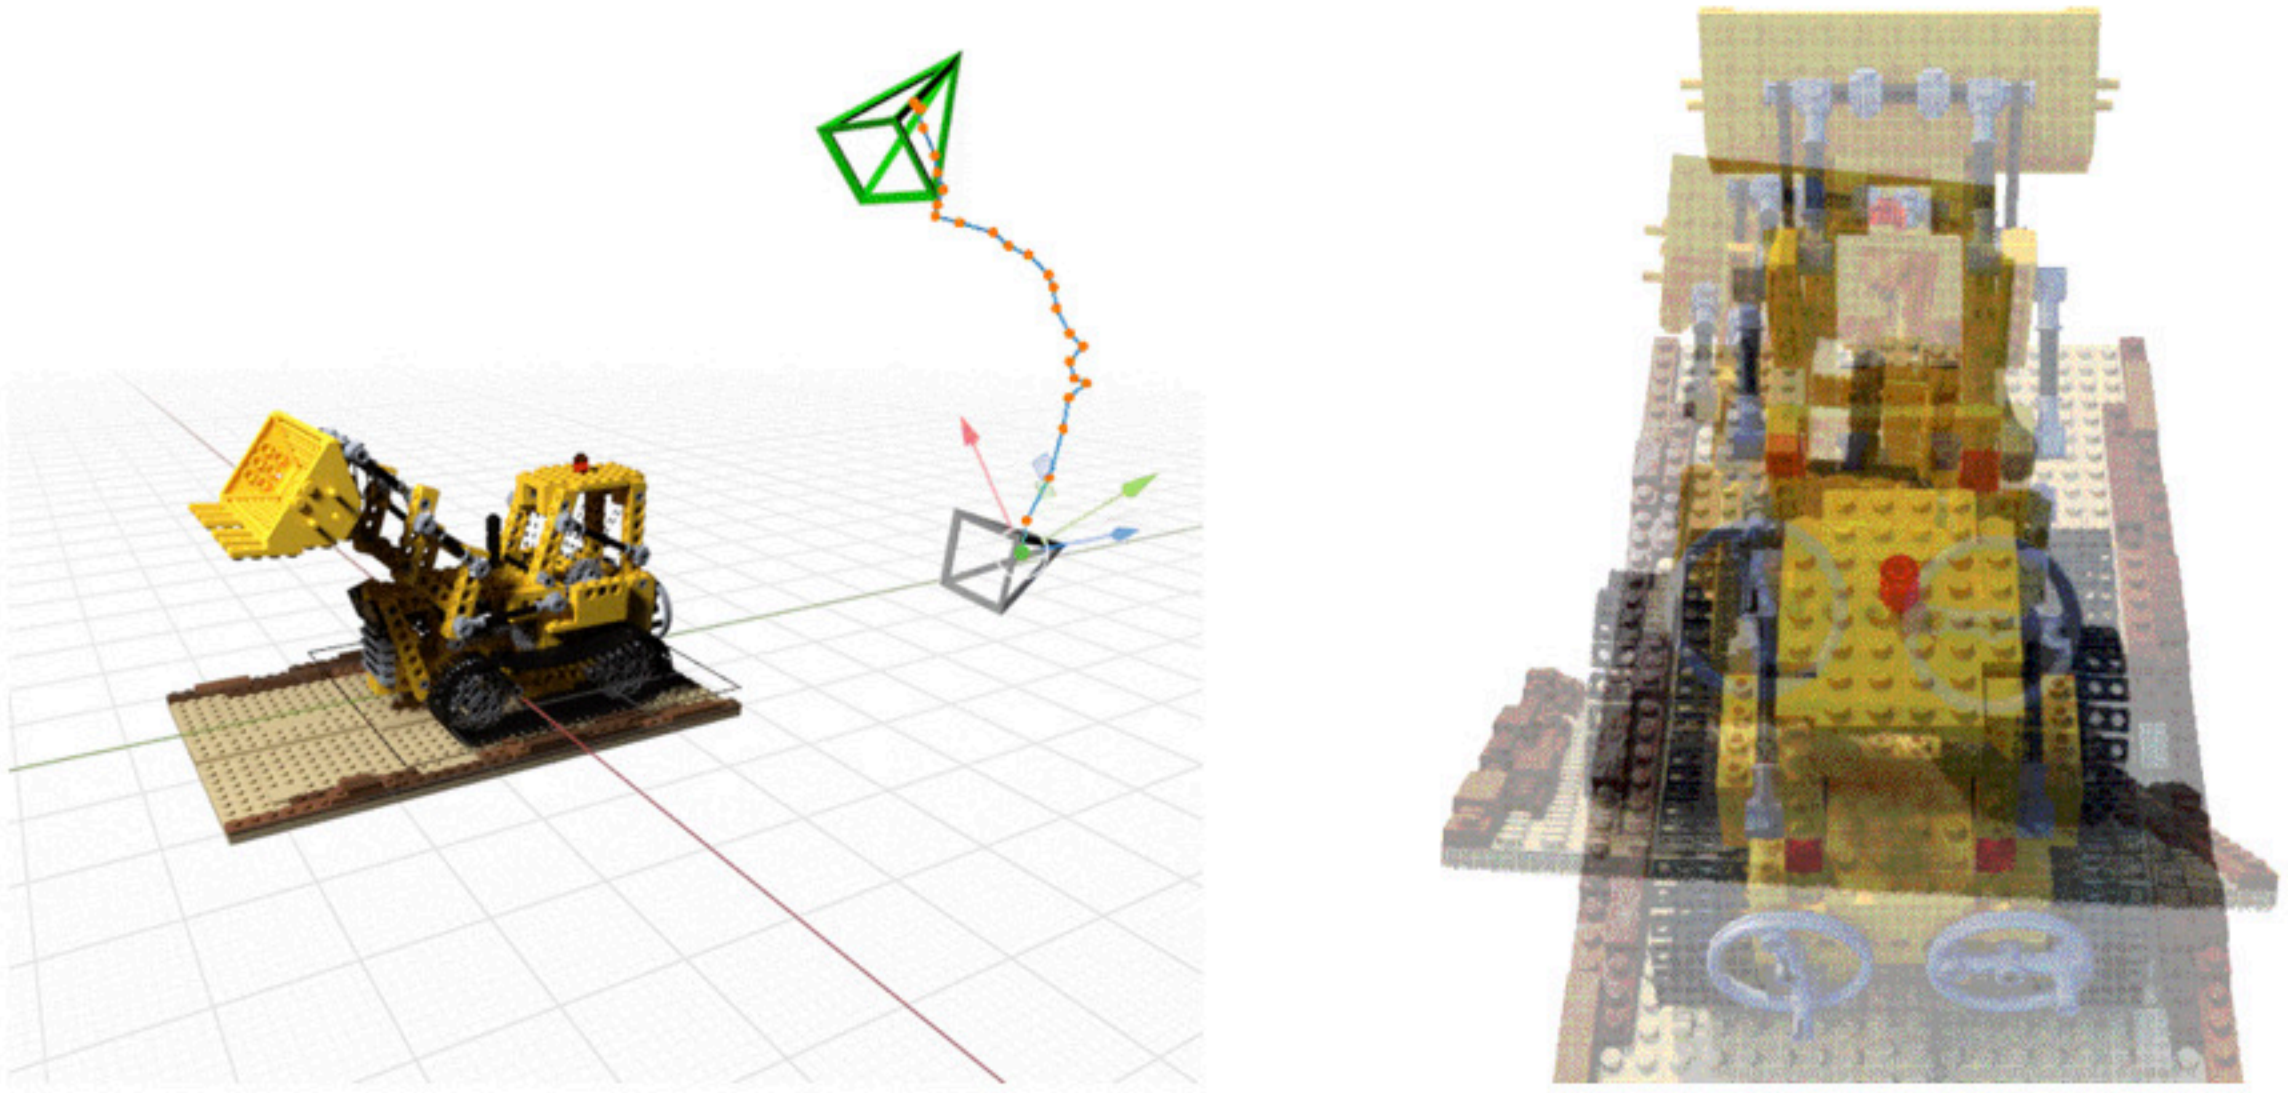

Yen-Chen et al. IROS'21
Inverting Neural Radiance Fields for Pose Estimation

nerf2nerf

Scene A

Scene B

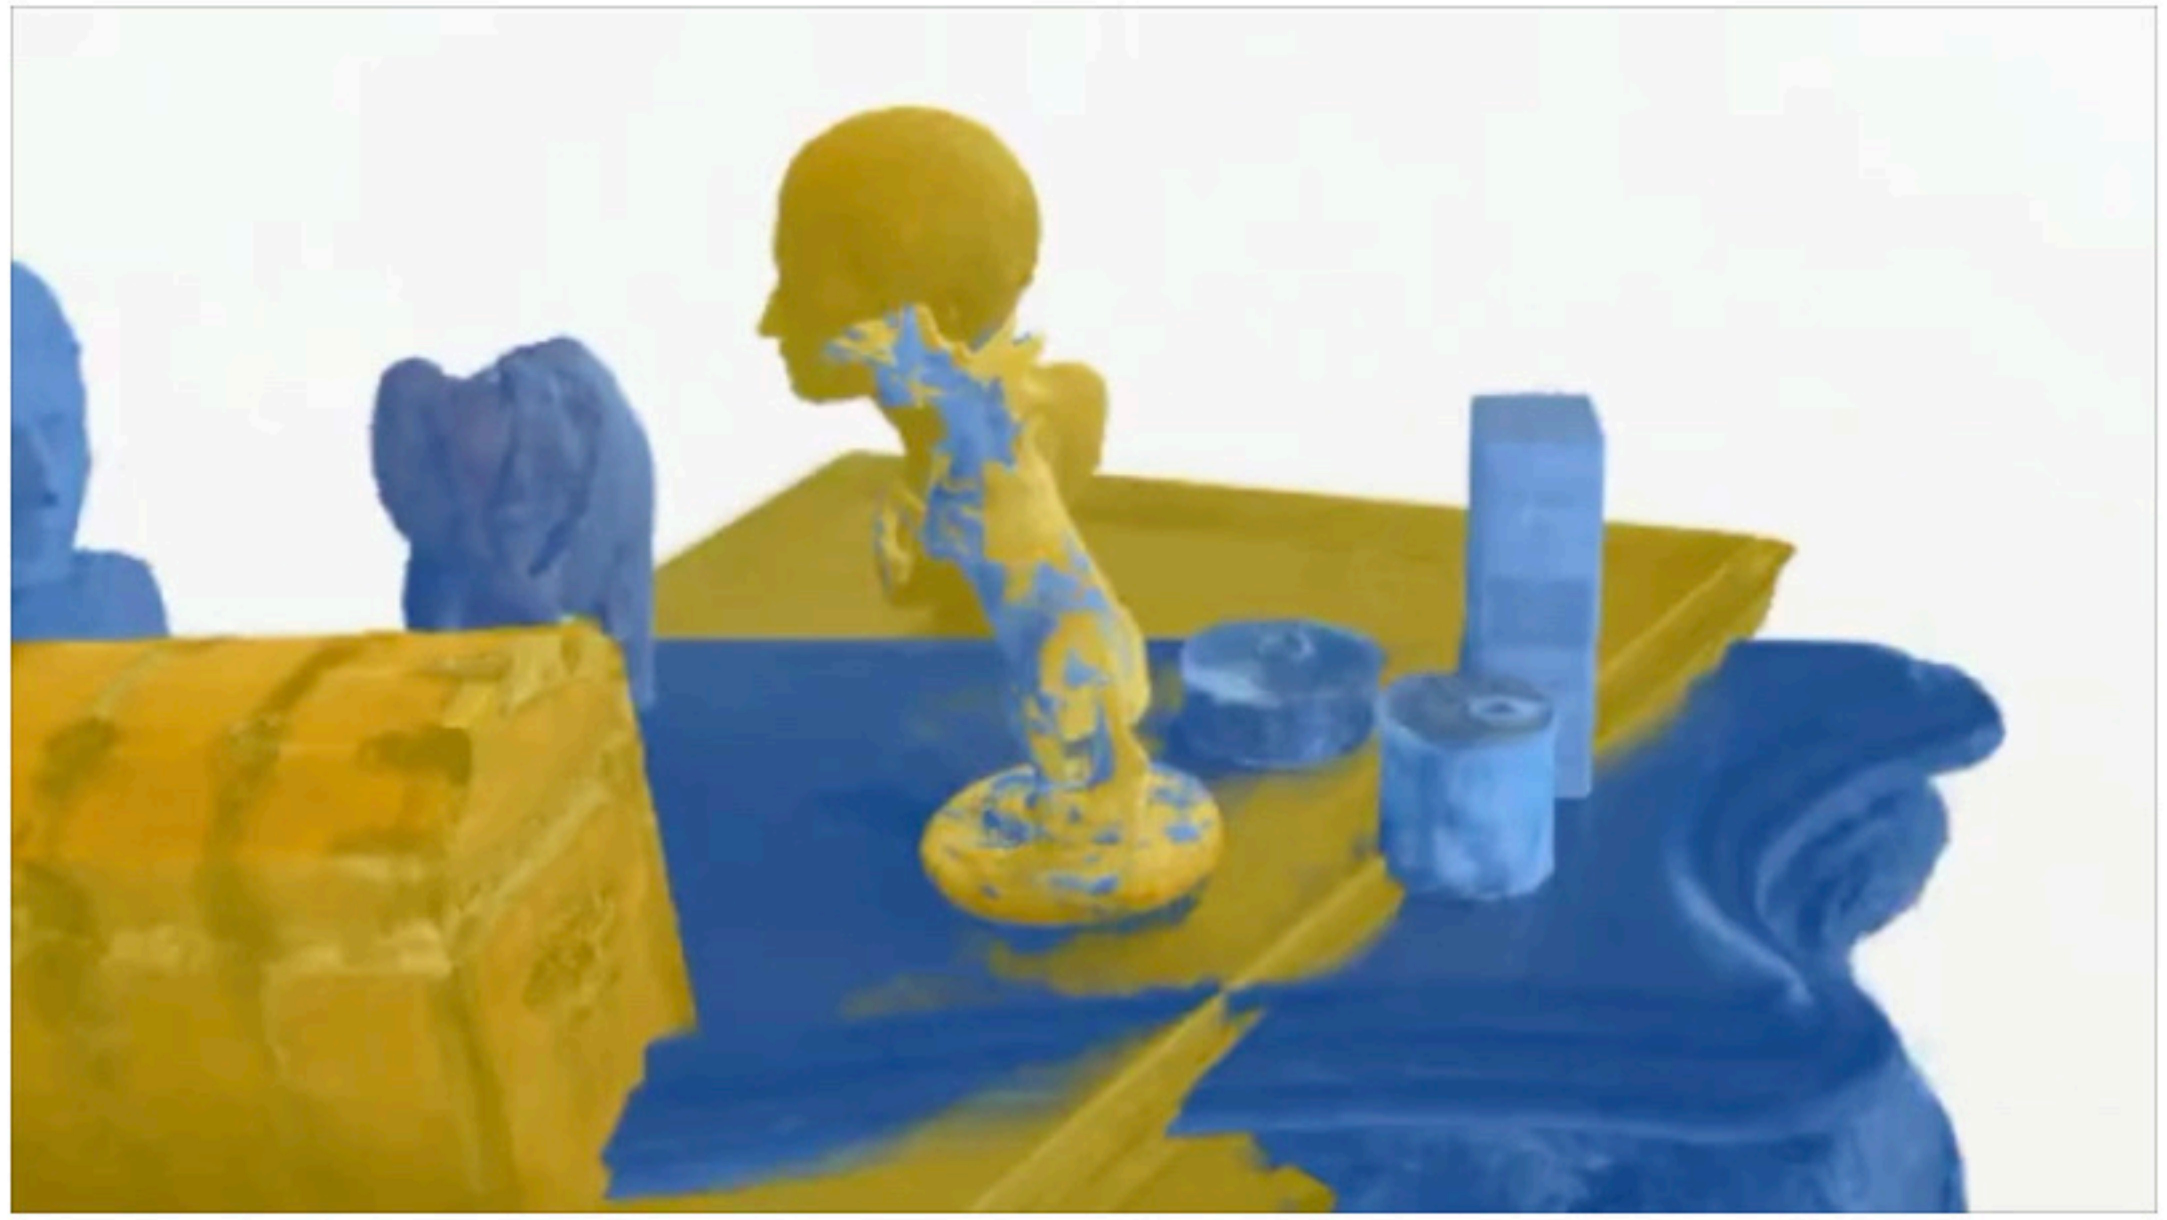

nerf2nerf

to register horse figurine

Lily Goli, Daniel Rebain, Sara Sabour, Animesh Garg, Andrea Tagliasacchi “nerf2nerf:Pairwise Registration of Neural Radiance Fields” ICRA 2023

Problem Setting

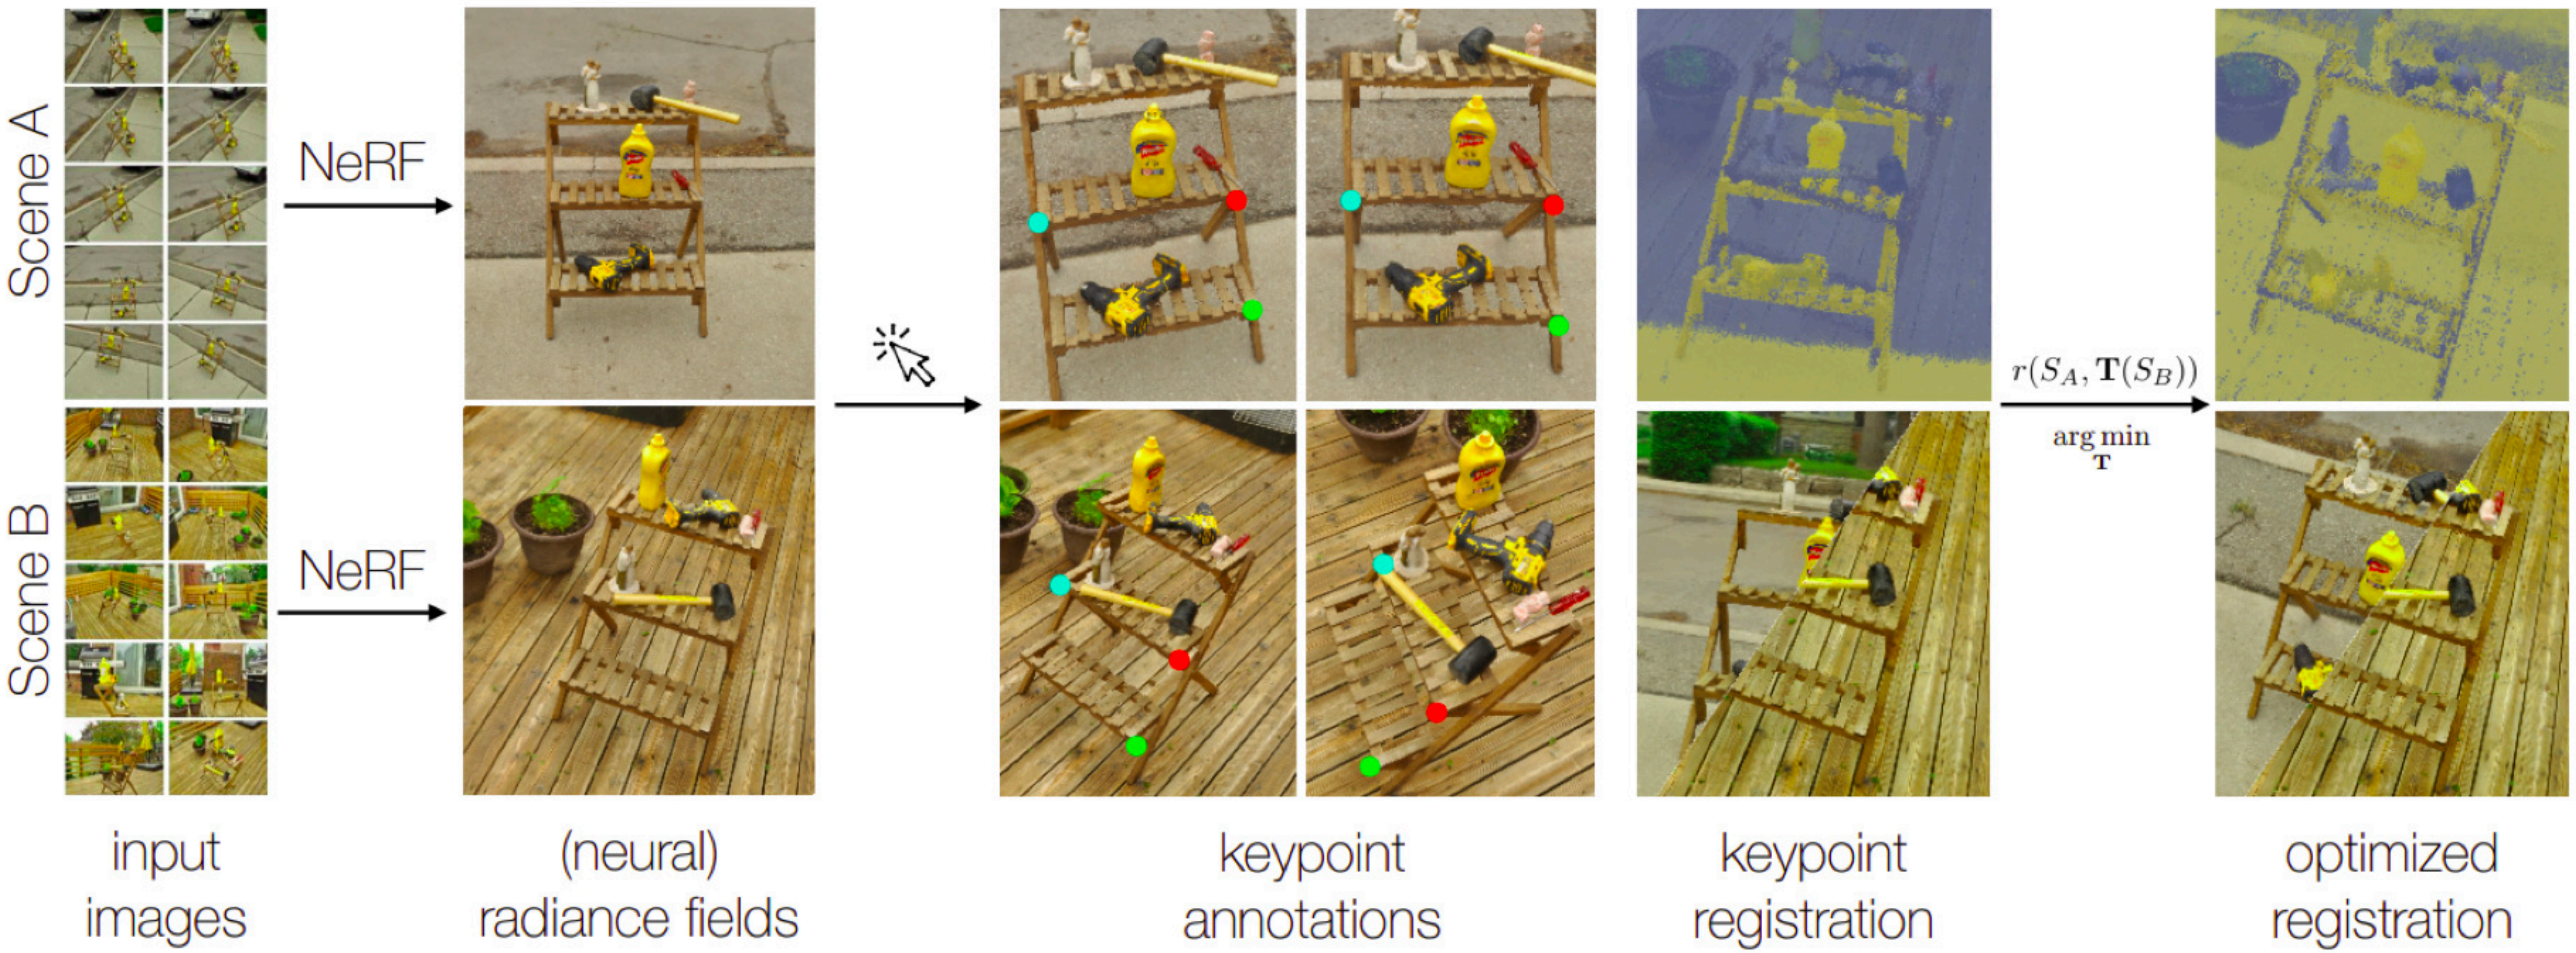

Talk Outline

How to extract a geometric representation from NeRF that can be used for the registration task?

- Should accurately capture object geometry
- Should be invariant to scene illumination
- Should be invariant to view direction


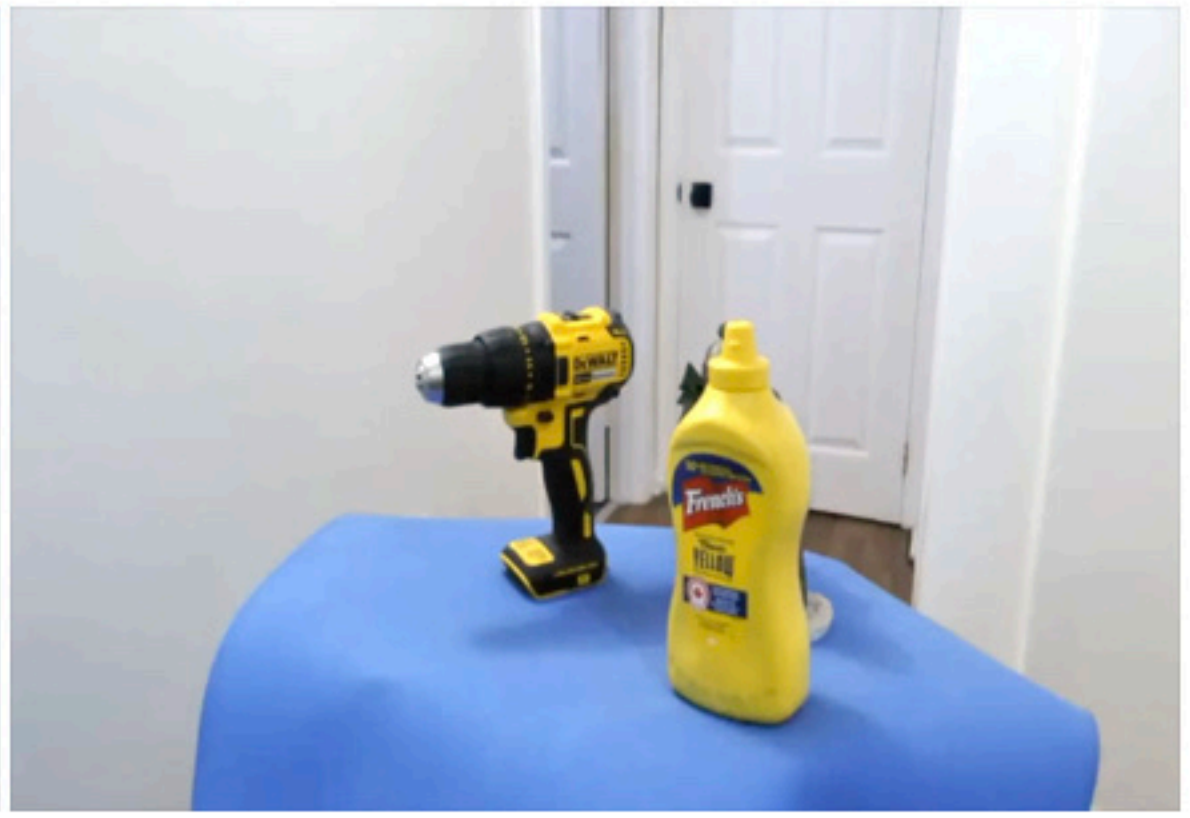

Lily Goli, Daniel Rebain, Sara Sabour, Animesh Garg, Andrea Tagliasacchi “nerf2nerf:Pairwise Registration of Neural Radiance Fields” ICRA 2023

Talk Outline

How to register two neural fields?

- work on Eulerian representation
- robust to partial overlap

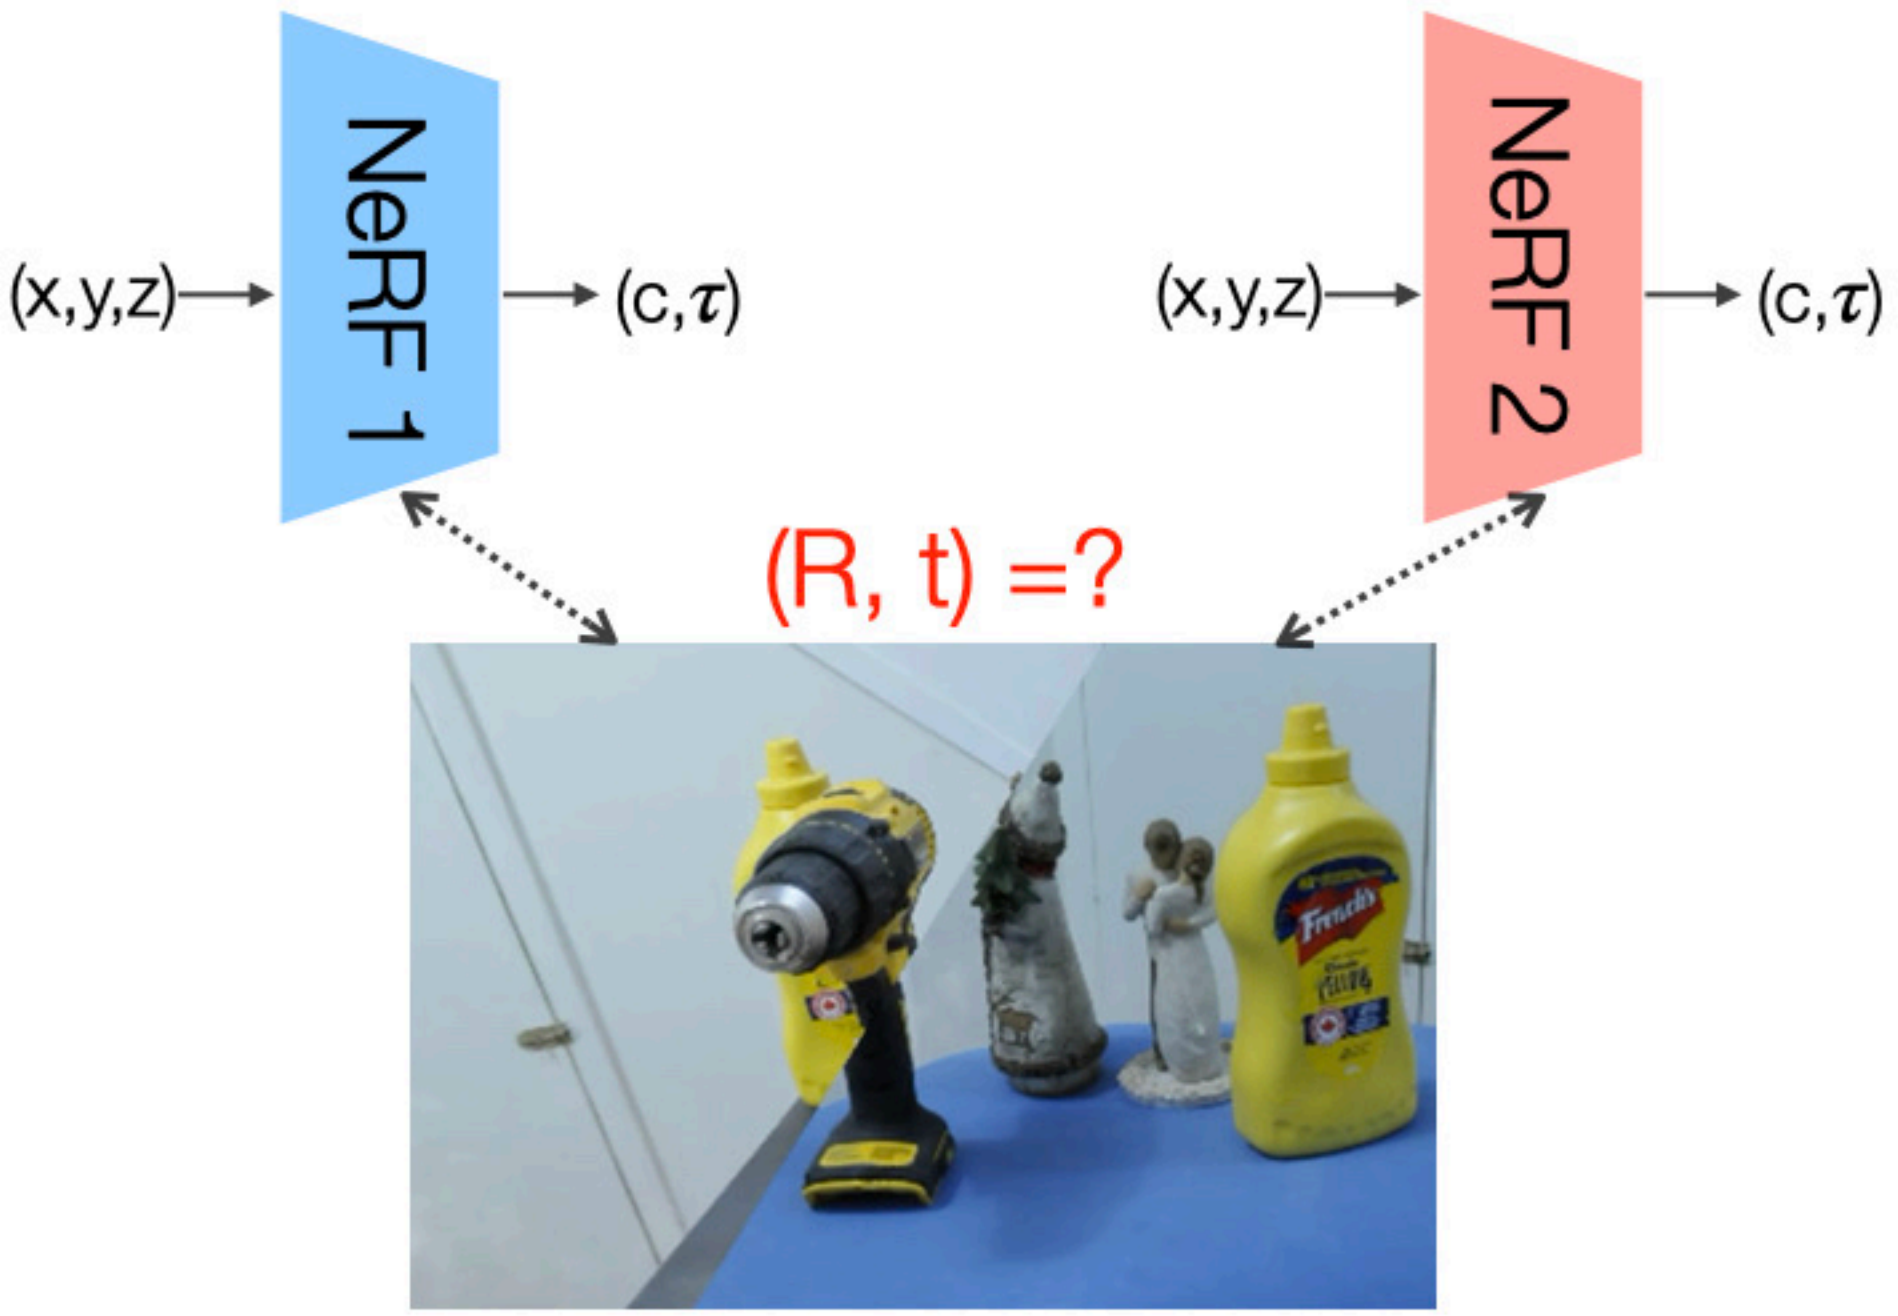

Surface Fields

Two products of NeRF:

Density $\tau(s)$: differential probability of hitting a solid particle at a point s.

Transmittance $\mathcal{T}(0 \rightarrow s \mid r)$: the probability that ray r hits no solid particles on its way to point s

NeRF

Density

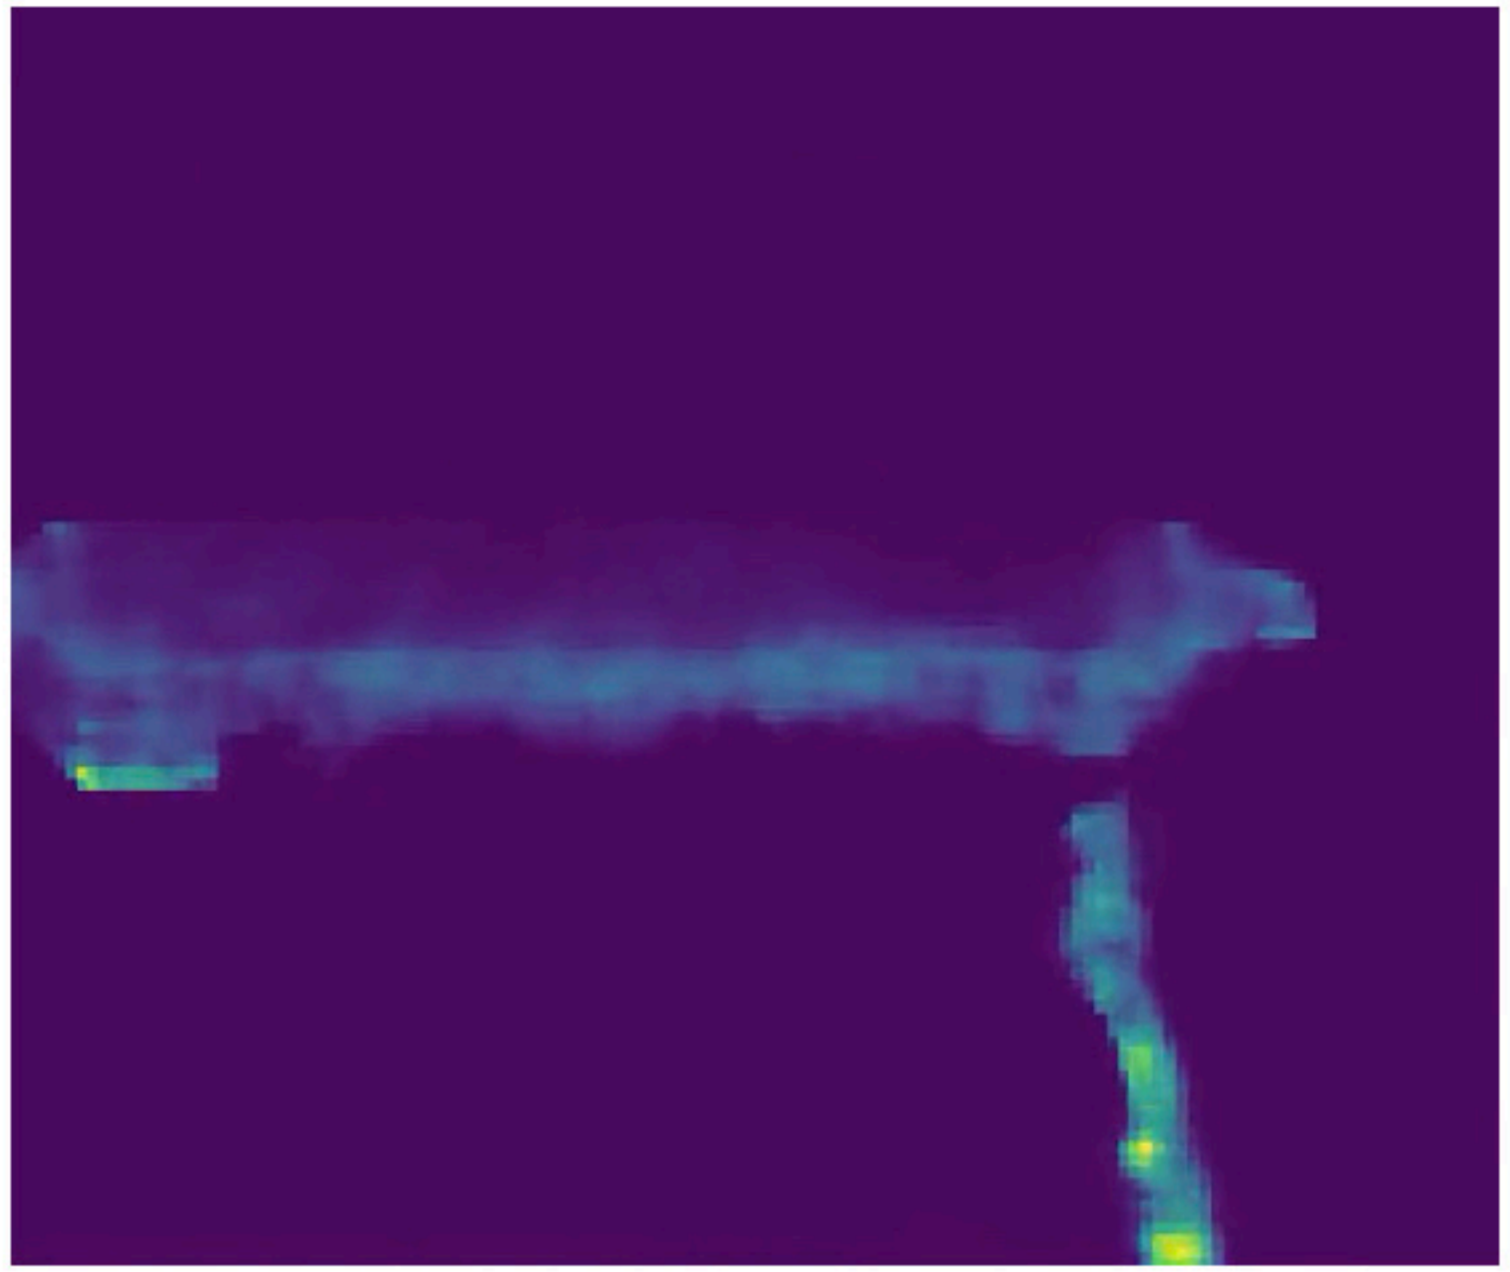

Transmittance

Lily Goli, Daniel Rebain, Sara Sabour, Animesh Garg, Andrea Tagliasacchi “nerf2nerf:Pairwise Registration of Neural Radiance Fields” ICRA 2023

Surface Fields

Density $\tau(s)$

Transmittance $\mathcal{T}(0 \rightarrow s \,|\, r)$

Surface

$$\mathcal{S}(s \,|\, r) = \int_{s-\delta}^{s+\delta} \mathcal{T}(0 \rightarrow t \,|\, r) \cdot \tau(t \,|\, r) \, dt$$

Density Transmittance Surface

Lily Goli, Daniel Rebain, Sara Sabour, Animesh Garg, Andrea Tagliasacchi “nerf2nerf:Pairwise Registration of Neural Radiance Fields” ICRA 2023

Surface Fields

Density $\tau(s)$

Transmittance $\mathscr{T}(0 \rightarrow s \mid r)$

Surface

$$\mathcal{S}(s \mid r) = \int_{s-\delta}^{s+\delta} \mathscr{T}(0 \rightarrow t \mid r) \cdot \tau(t \mid r)\, dt$$

$$\mathcal{S}(s) = \max_{\mathbf{o} \in \mathcal{O}} \mathcal{S}\left(\|\mathbf{o} - s\| \mid \left(\mathbf{o}, \frac{\mathbf{o} - s}{\|\mathbf{o} - s\|} \right) \right) \in [0, 1]$$

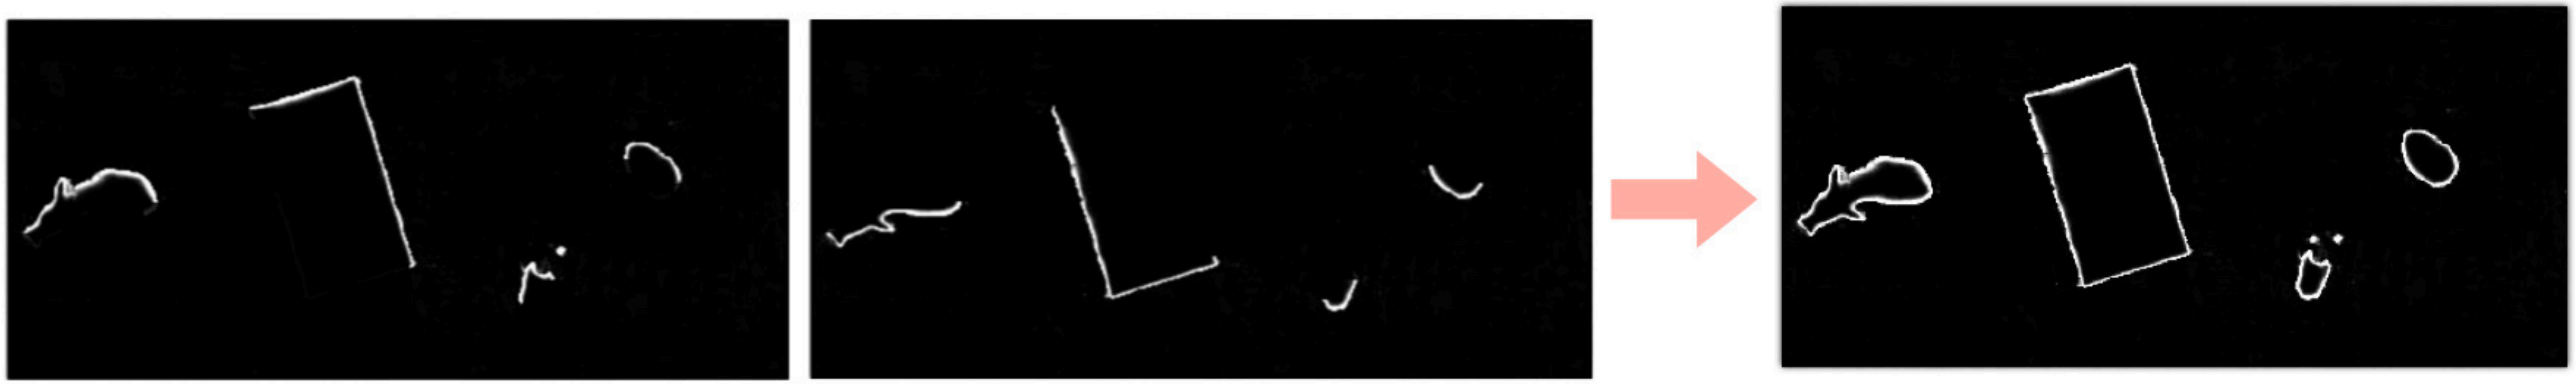

Surface Field

To have a conservative estimate of the field we can threshold it:

$$\mathcal{S}^{\epsilon}(\mathbf{x}) = \mathbb{1}(\mathcal{S}(\mathbf{x}) > \epsilon) \in \{0, 1\}$$

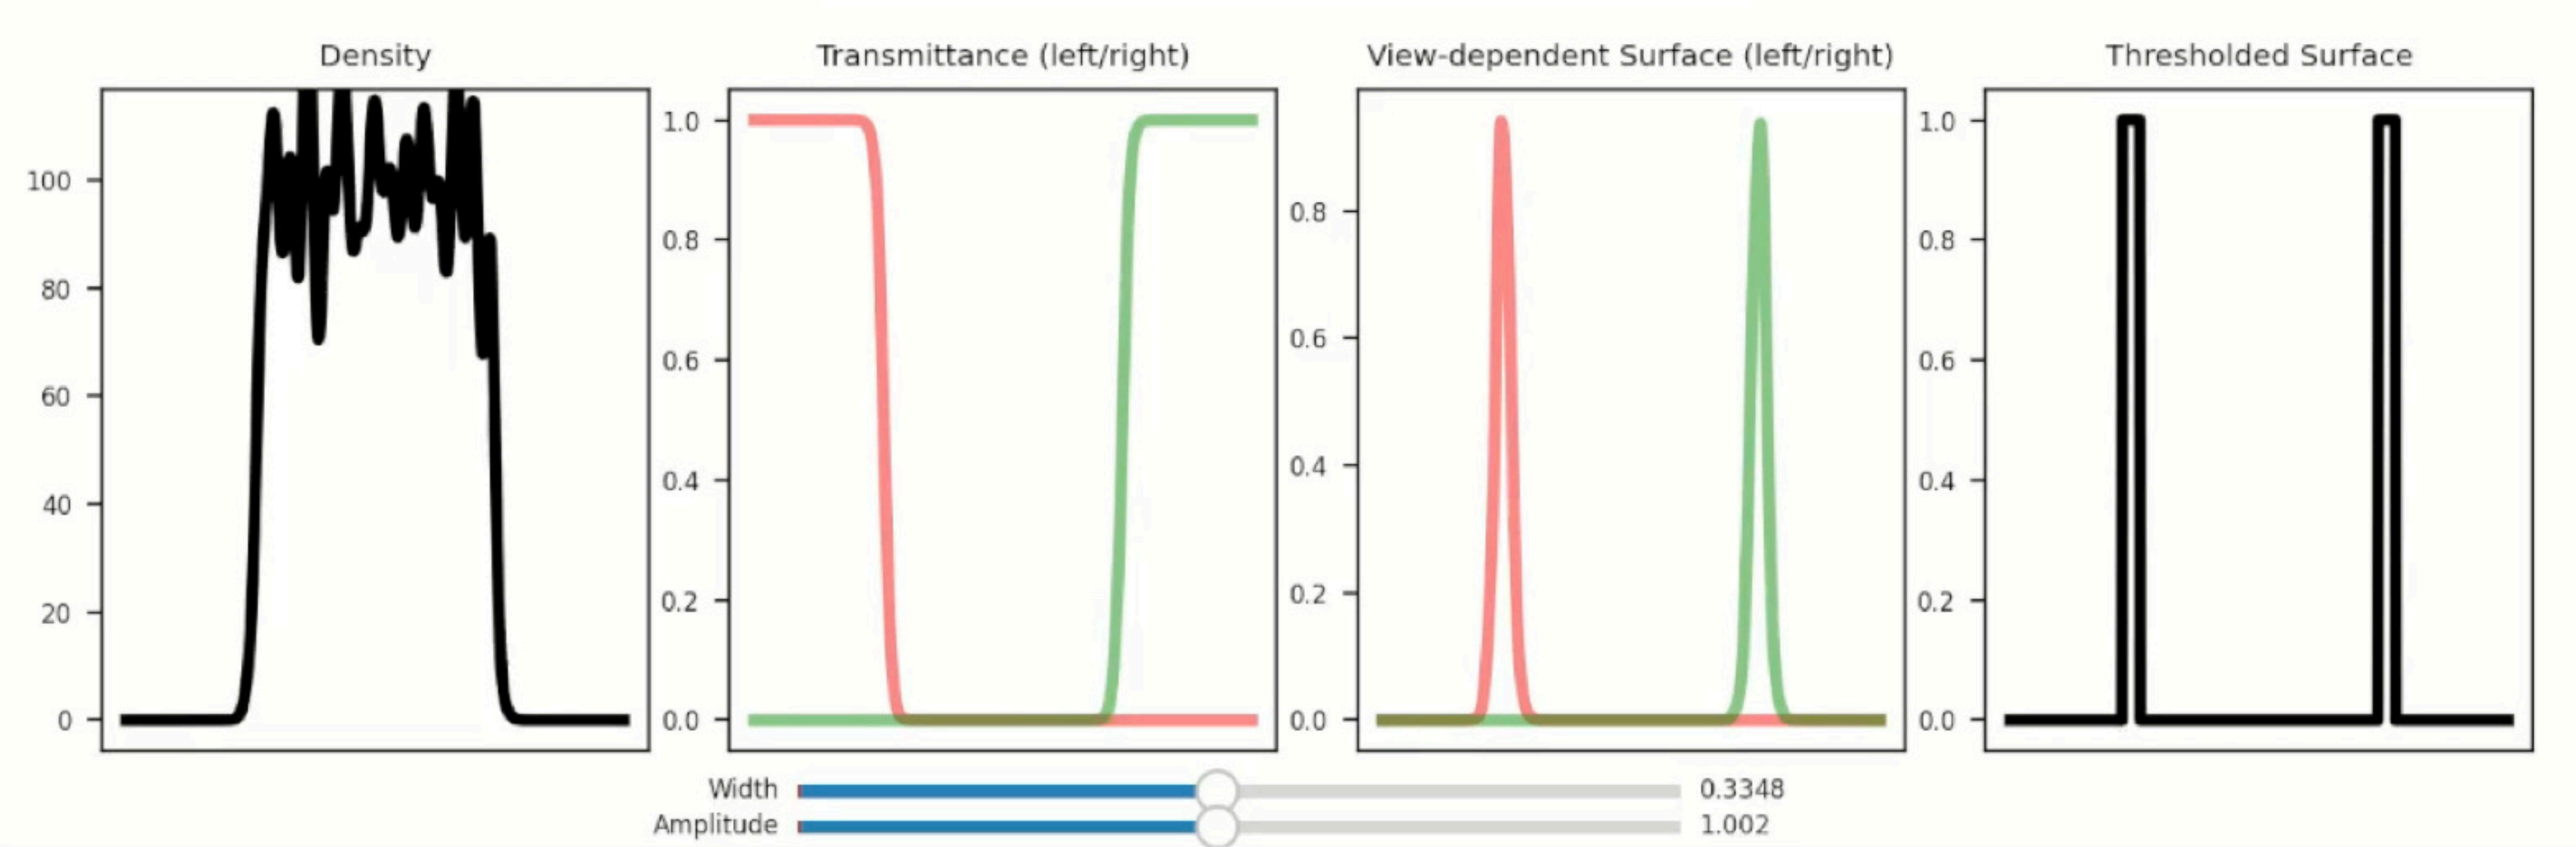

Goli et al. “nerf2nerf:Pairwise Registration of Neural Radiance Fields” ICRA 2023

Surface Field

To have a conservative estimate of the field we can threshold it:

$$\mathcal{S}^{\epsilon}(\mathbf{x}) = \mathbb{1}(\mathcal{S}(\mathbf{x}) > \epsilon) \in \{0, 1\}$$

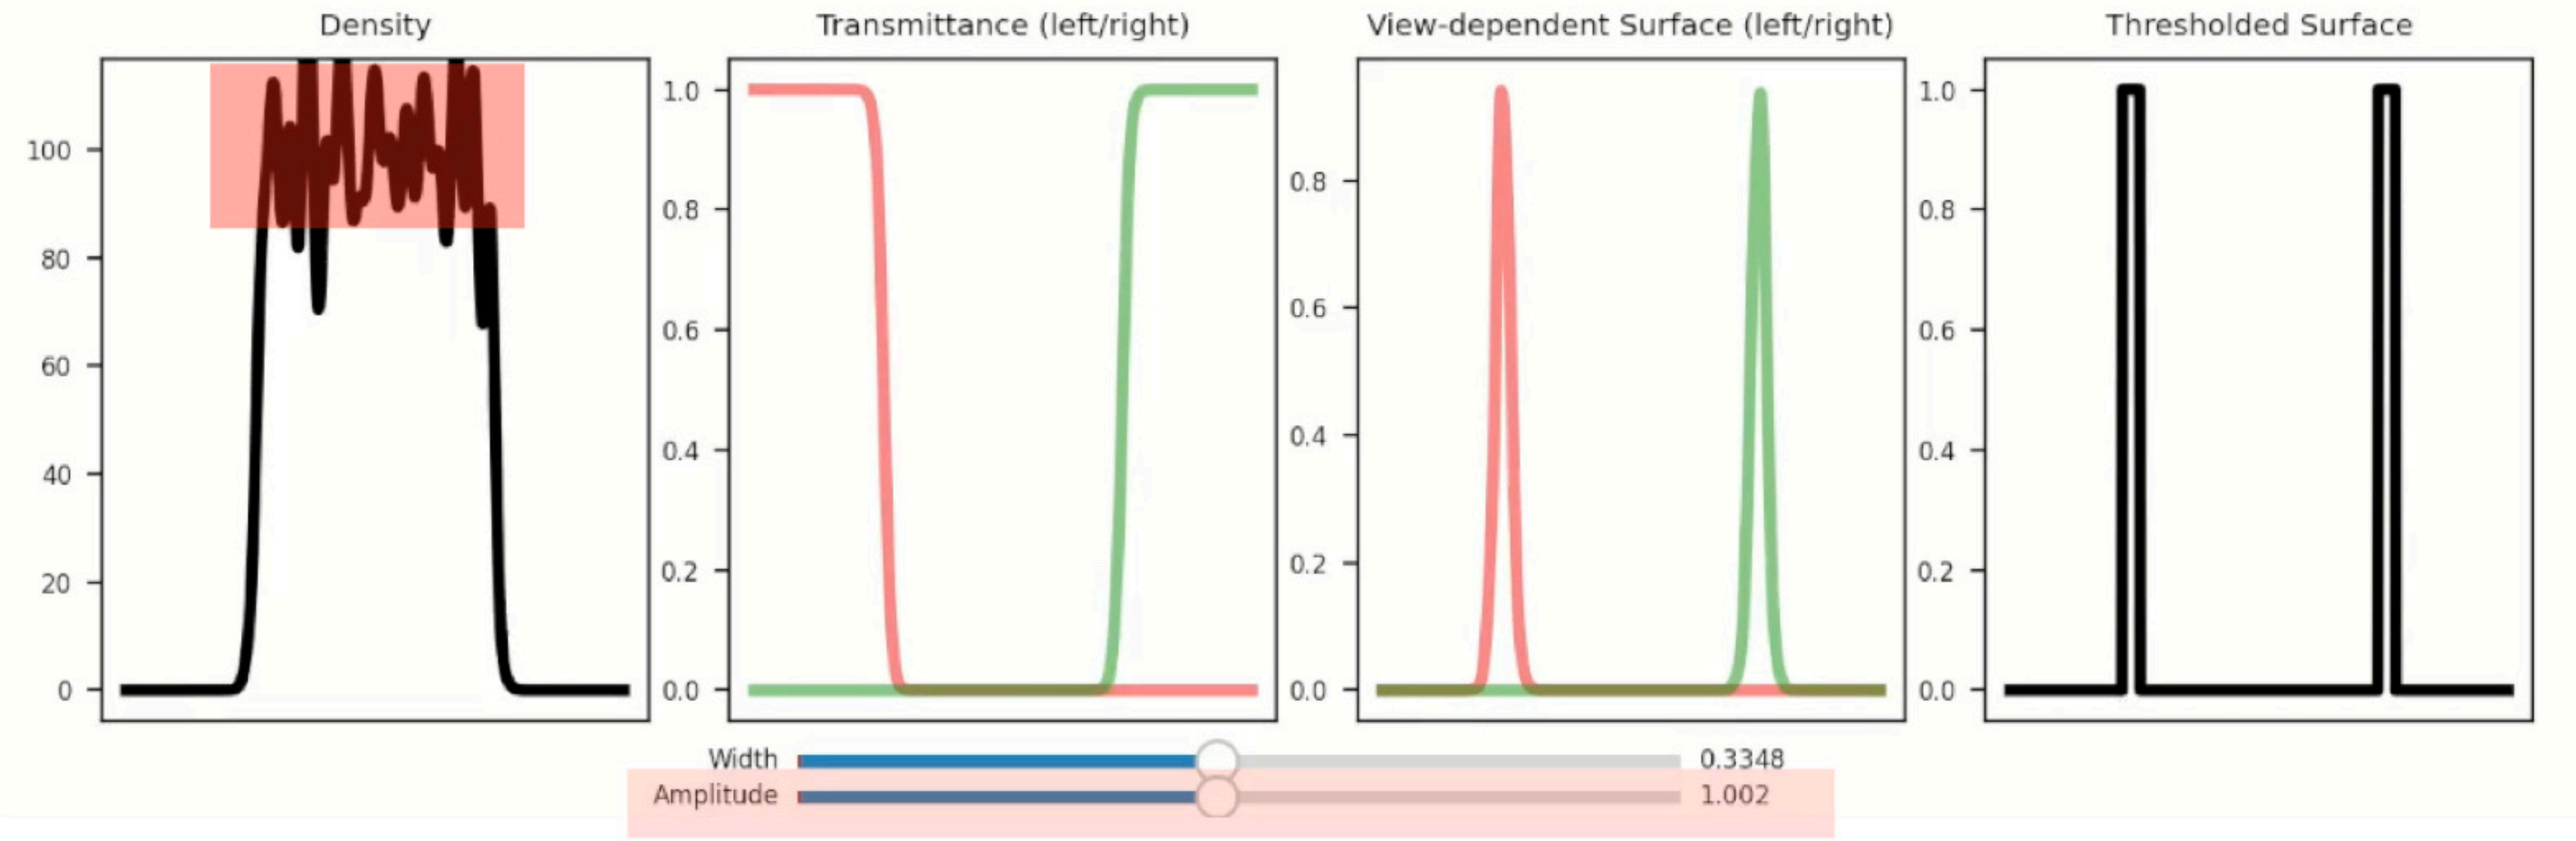

Goli et al. “nerf2nerf:Pairwise Registration of Neural Radiance Fields” ICRA 2023

Surface Field

To have a conservative estimate of the field we can threshold it:

$$\mathcal{S}^{\epsilon}(\mathbf{x}) = \mathbb{1}(\mathcal{S}(\mathbf{x}) > \epsilon) \in \{0, 1\}$$

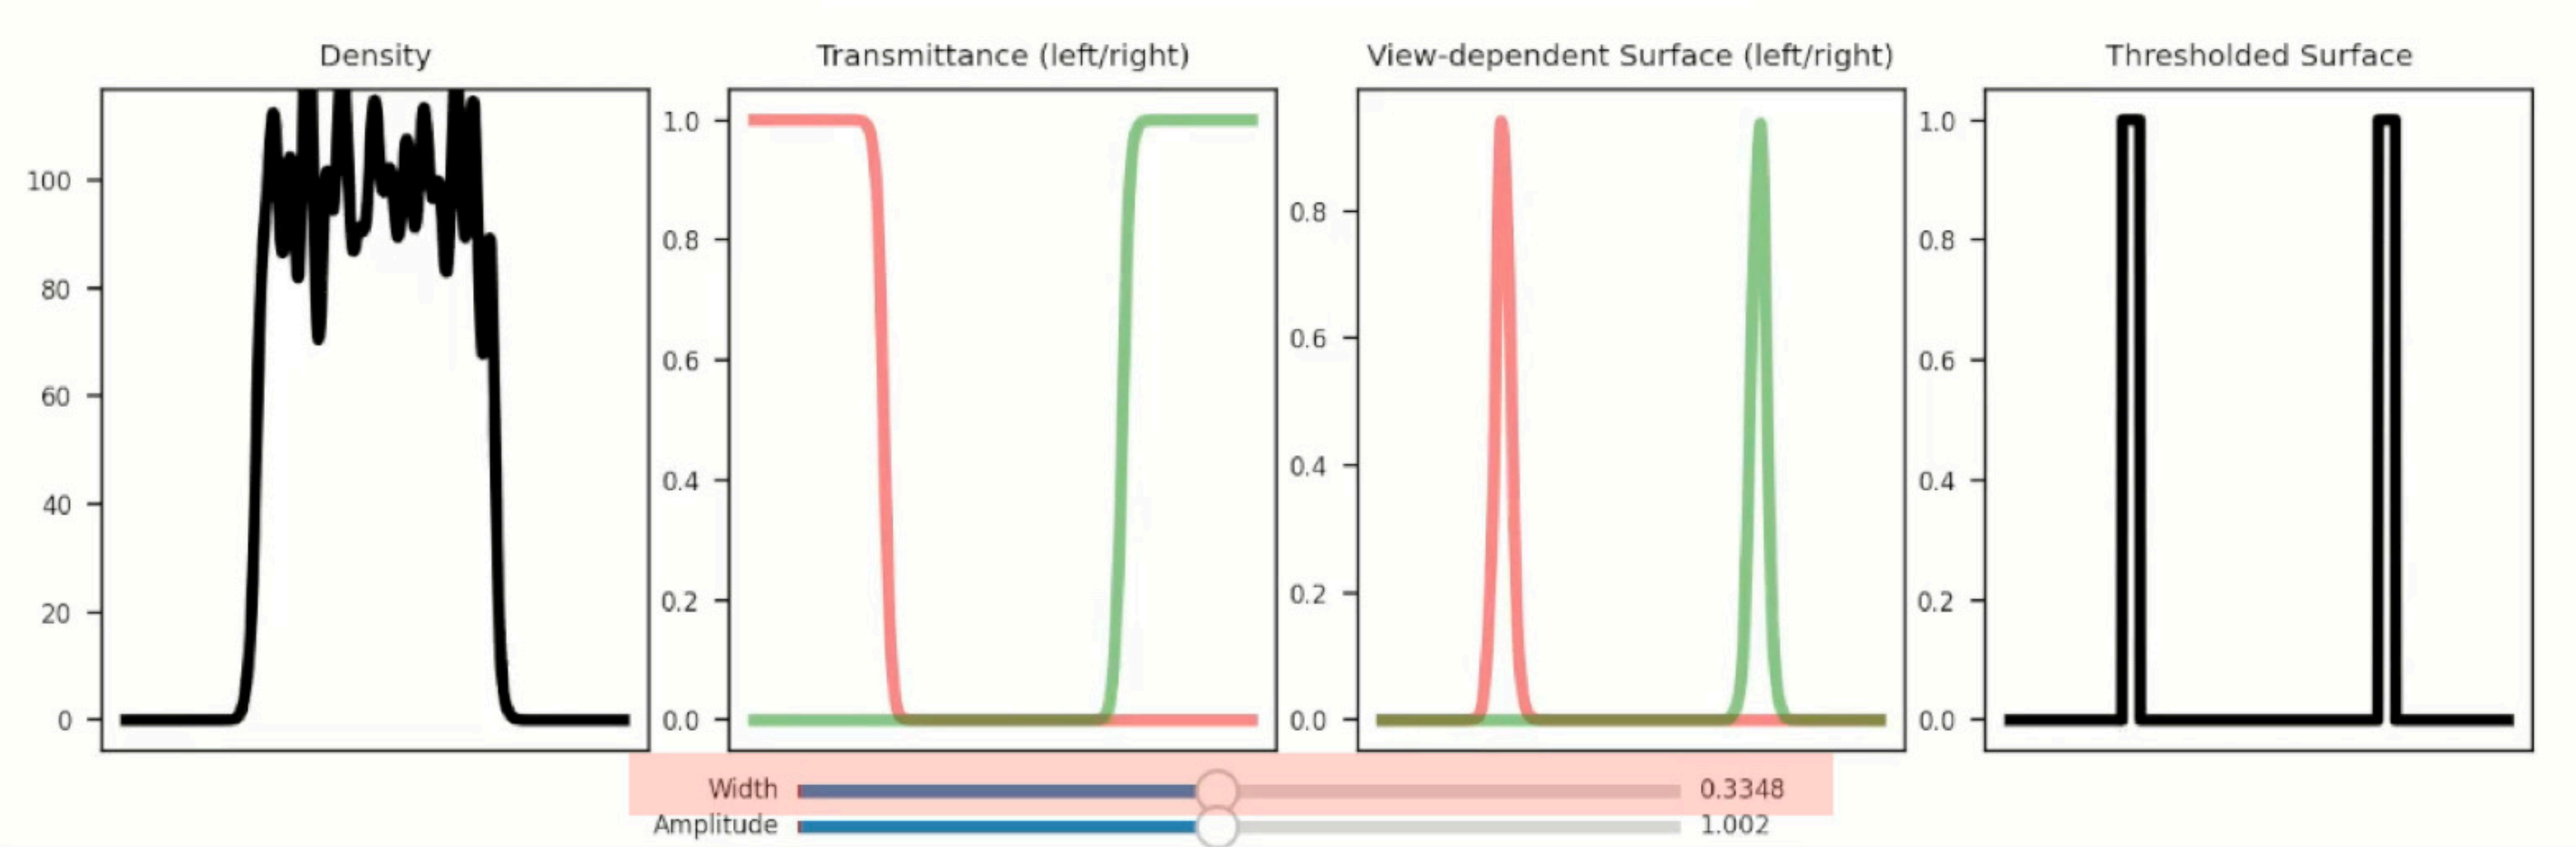

Distilling Surface Fields

The thresholded surface has a categorical co-domain an is not suitable for differentiation.

We can smooth it with a Gaussian convolution:

$$\mathcal{S}^{\sigma}(\mathbf{x}) = \mathcal{S}^{\epsilon}(\mathbf{x}) \circledast \mathcal{N}(0, \sigma^2) \approx \mathbb{E}_{\mathbf{z}\sim\mathcal{N}(\mathbf{x},\sigma^2)} \; [\mathcal{S}^{\epsilon}(\mathbf{z})]$$

To have faster query time on surface field we can distill it into another network.

We use integrated positional encoding[1] to encode different levels of smoothness:

$$\gamma_{\sigma}(\mathbf{x}) = \mathbb{E}_{\mathbf{z}\in\mathcal{N}(\mathbf{x},\sigma^2)}[\gamma(\mathbf{z})] = \left\{ \left(\frac{sin(2^l \mathbf{z})\cdot}{\sqrt{exp(4^l \sigma^2)}}, \frac{cos(2^l \mathbf{z})\cdot}{\sqrt{exp(4^l \sigma^2)}} \right) \right\}_{l=0}^{L}$$

$$\arg\min_{\theta} \; \mathbb{E}_{\mathbf{x}\in\mathcal{B}(\mathbf{0},r)} \left[\mathcal{L}_{\text{poisson}}(\mathcal{S}^{\sigma}(\gamma_{\sigma}(\mathbf{x});\theta)), \; \mathcal{S}^{\sigma}(\mathbf{x})) \right]$$

$$\mathcal{L}_{\text{poisson}}(x, y) = x - y \log x$$

1. Barron et al “Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields,” in ICCV, 2021.

Distilling Surface Field

Scene

Surface $\sigma = 0$

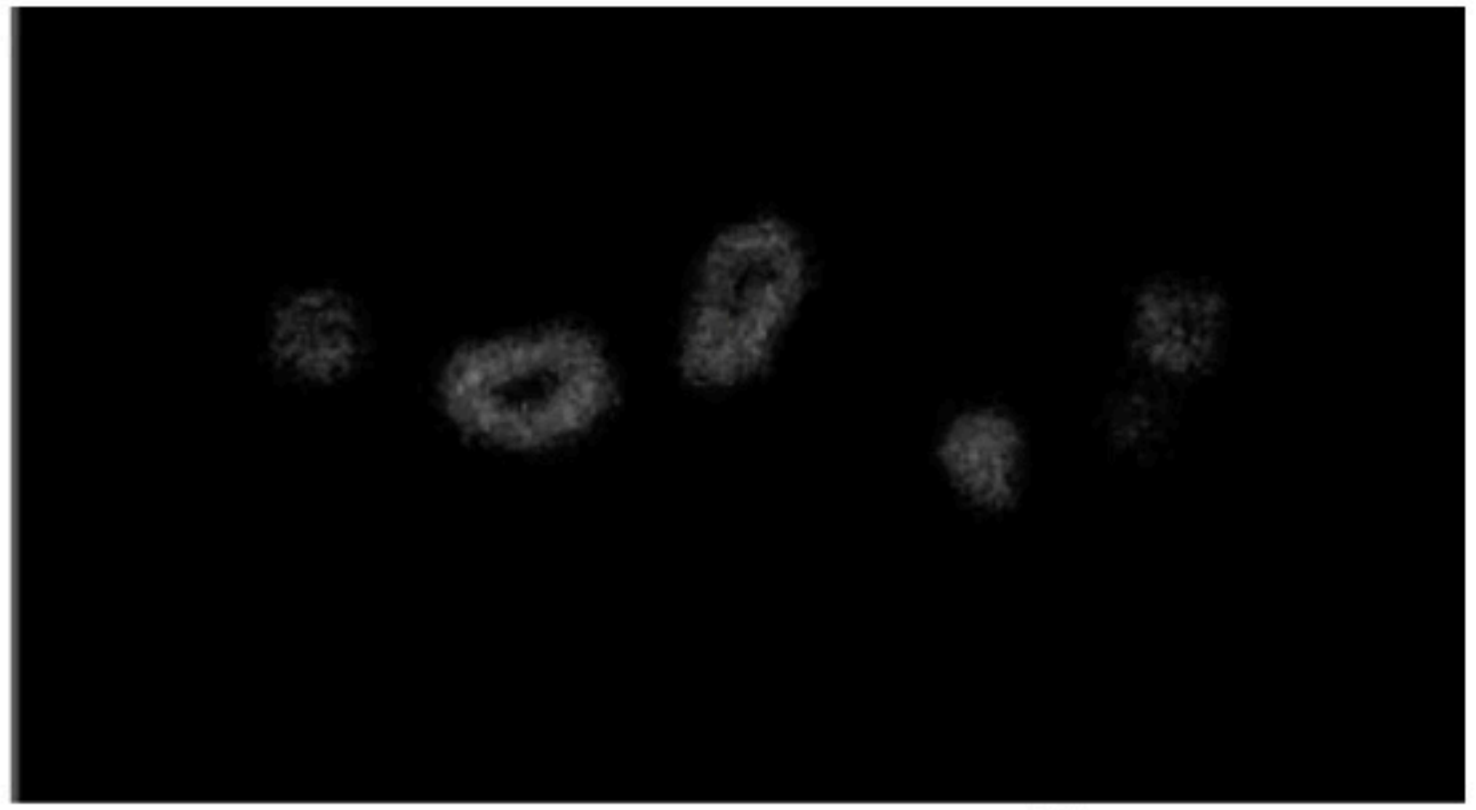

Surface $\sigma > 0$

Goli et al. “nerf2nerf:Pairwise Registration of Neural Radiance Fields” ICRA 2023

Talk Outline

How to extract a geometric representation from NeRF that can be used for the registration task?

- Should be invariant to scene illumination
- Should be invariant to view direction
- Should accurately capture object geometry

How to register two neural fields?

- robust to partial overlap — Optimization
- work on Eulerian representation — Sampling

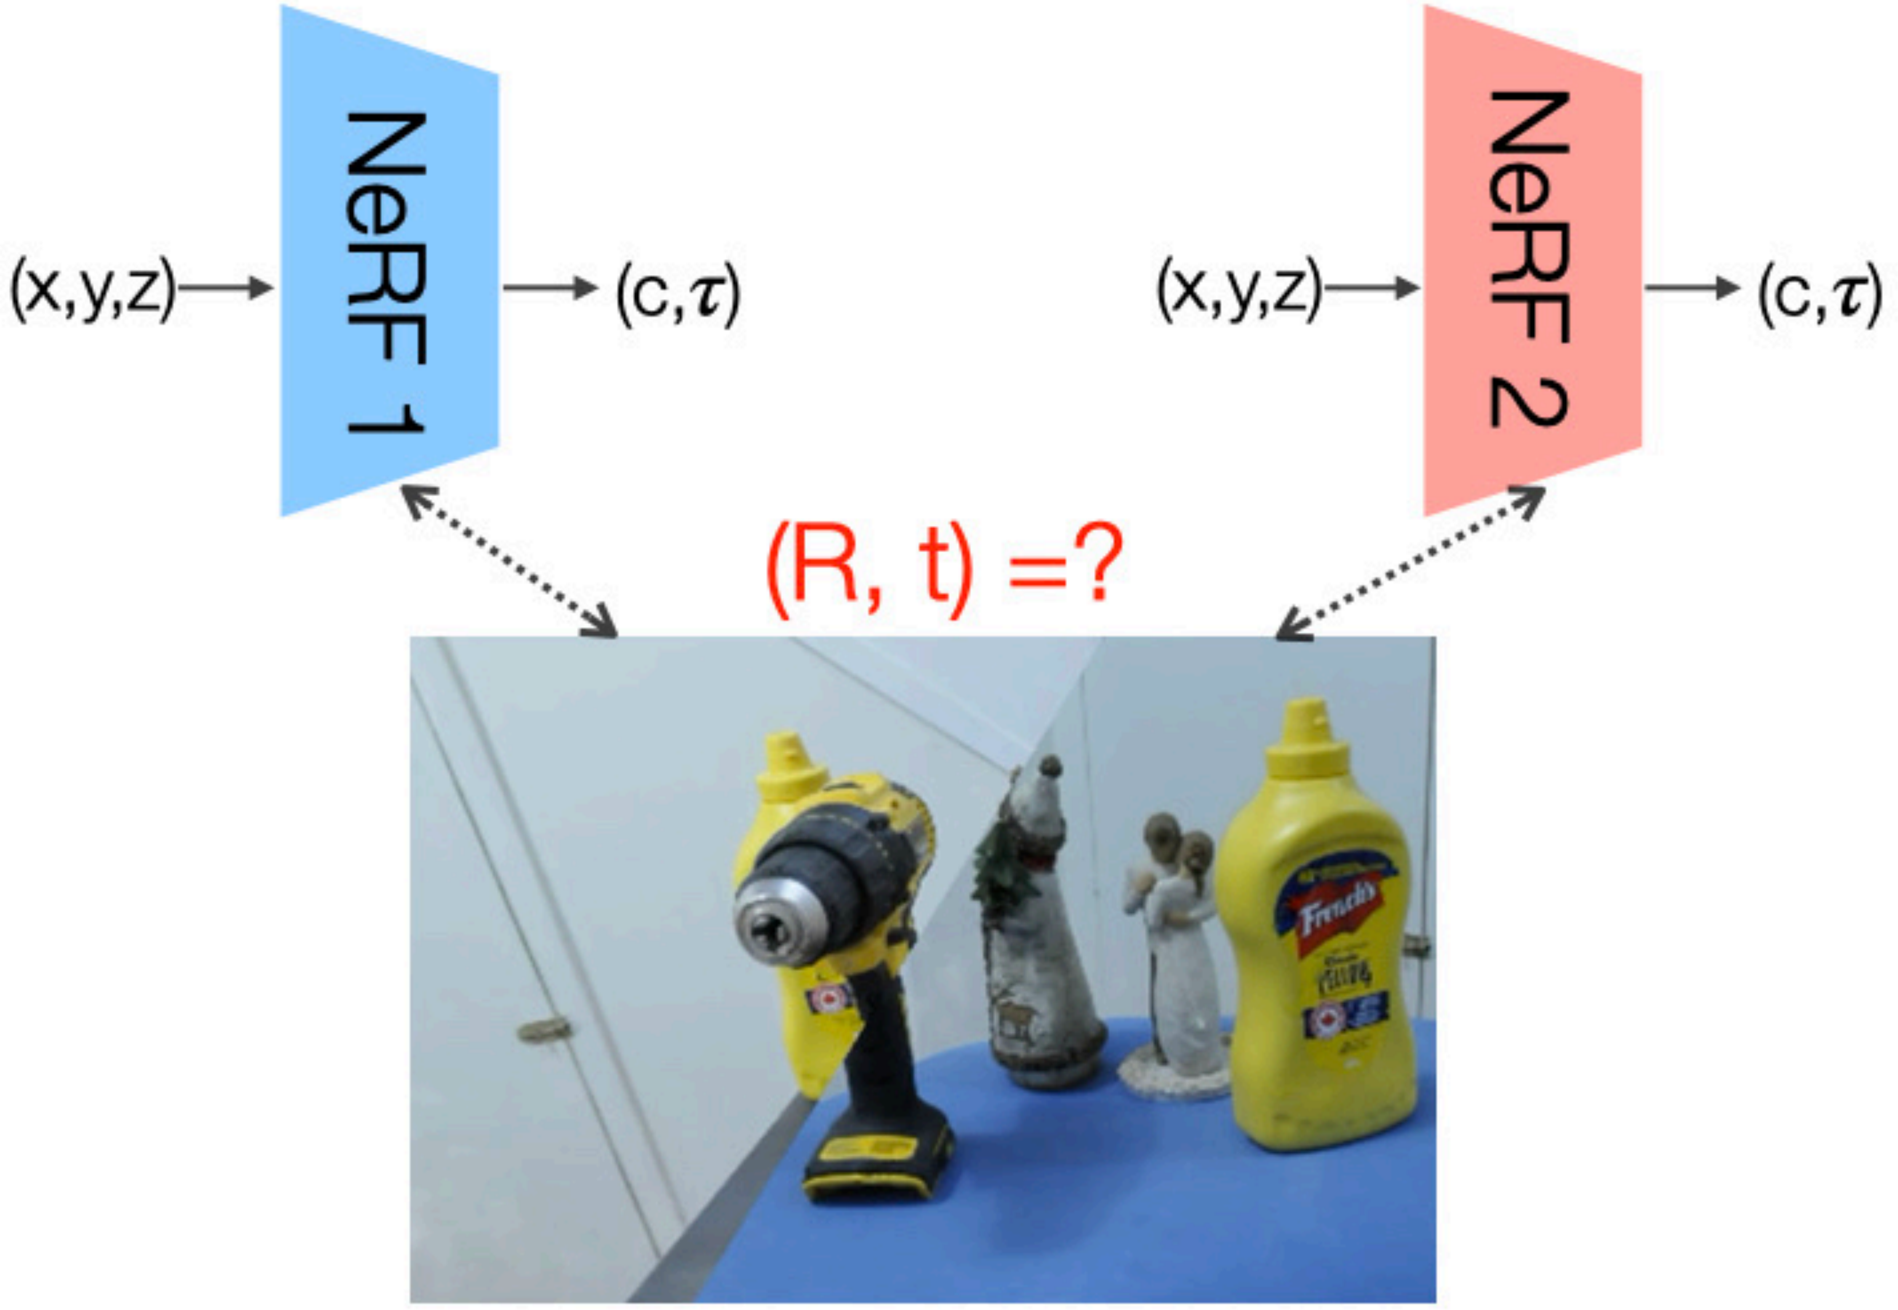

Energy-based Optimization

$$\boxed{\text{argmin}_{R,t}\mathcal{L}_{match} + \lambda\mathcal{L}_{key}}$$

$$r(x; \mathcal{S}_a^\sigma, \mathcal{S}_b^\sigma, R, t) = ||\mathcal{S}_a^\sigma(x) - \mathcal{S}_b^\sigma(Rx + t)||_2$$

$$\mathcal{L}_{match}(\mathcal{S}_a^\sigma, \mathcal{S}_b^\sigma, R, t) = \mathbb{E}_{x \in \mathcal{A}}\kappa(||\mathcal{S}_a^\sigma(x) - \mathcal{S}_b^\sigma(Rx + t)||_2; c, \alpha)$$

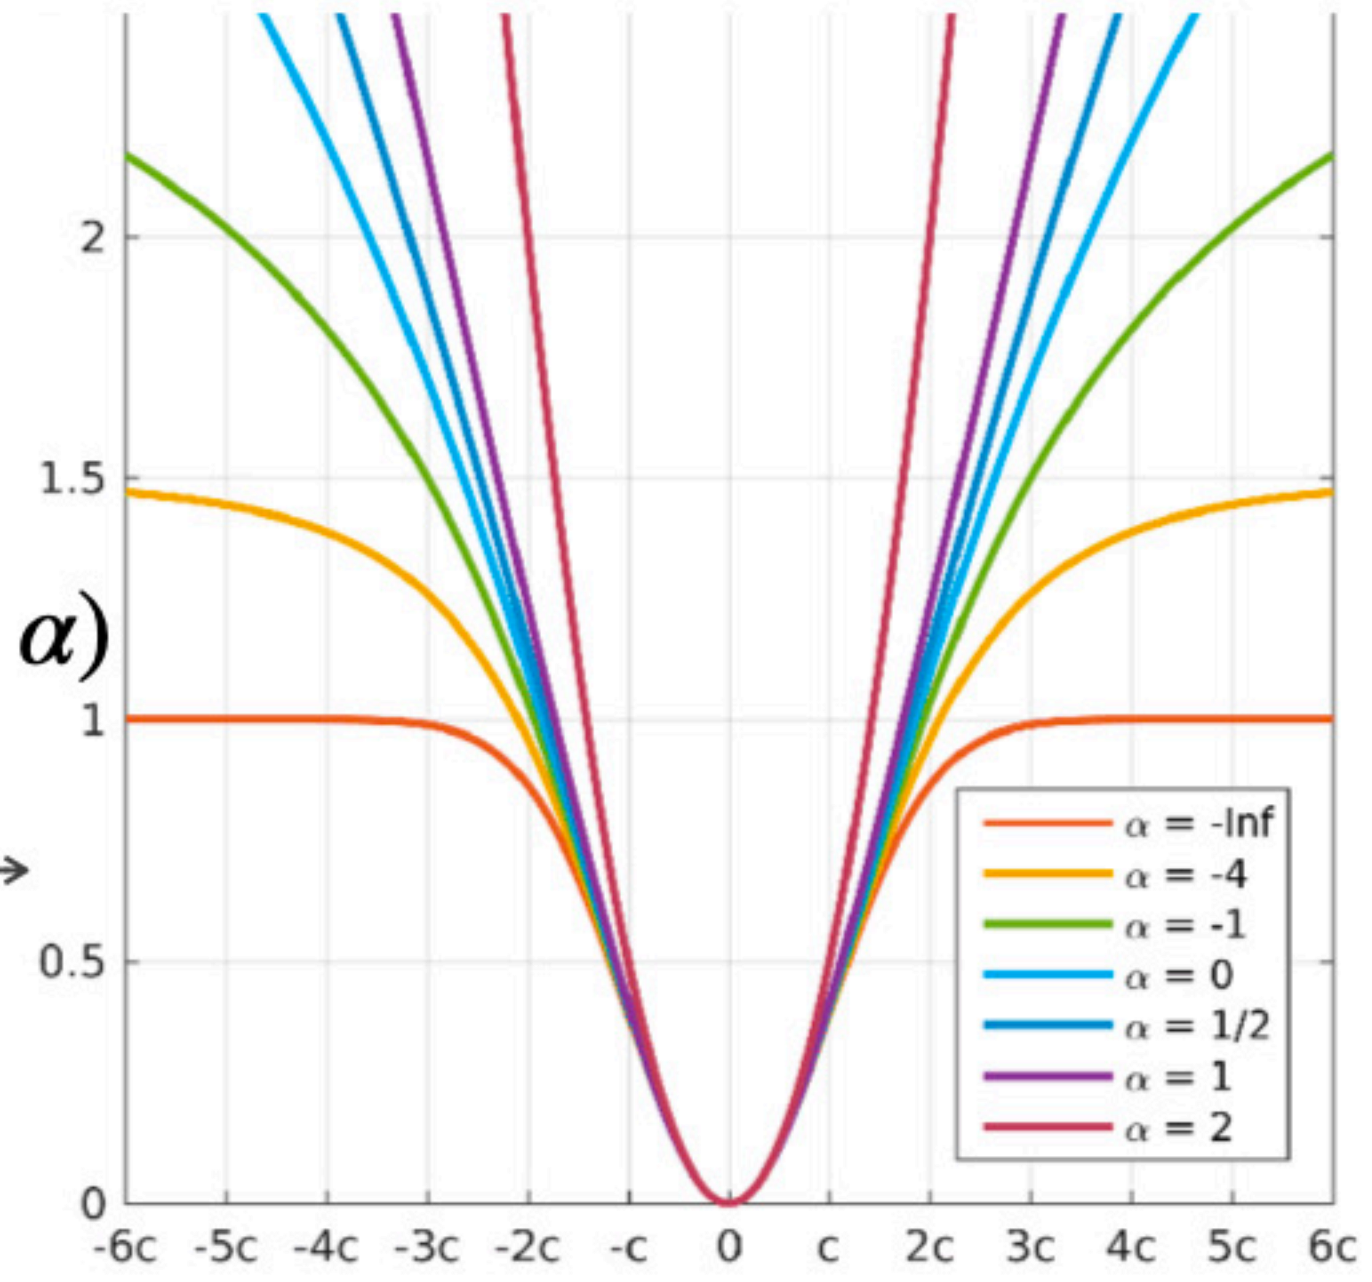

$$\mathcal{L}_{key}(R, t; \mathcal{Q}_a, \mathcal{Q}_b) = \sum_{q_a, q_b} ||q_a - (Rq_b + t)||_2^2$$

Sampling

Metropolis-Hastings Sampling: to sample points near key-point locations that

- are near surface
- are far enough and informative
- are in correspondence according to residual

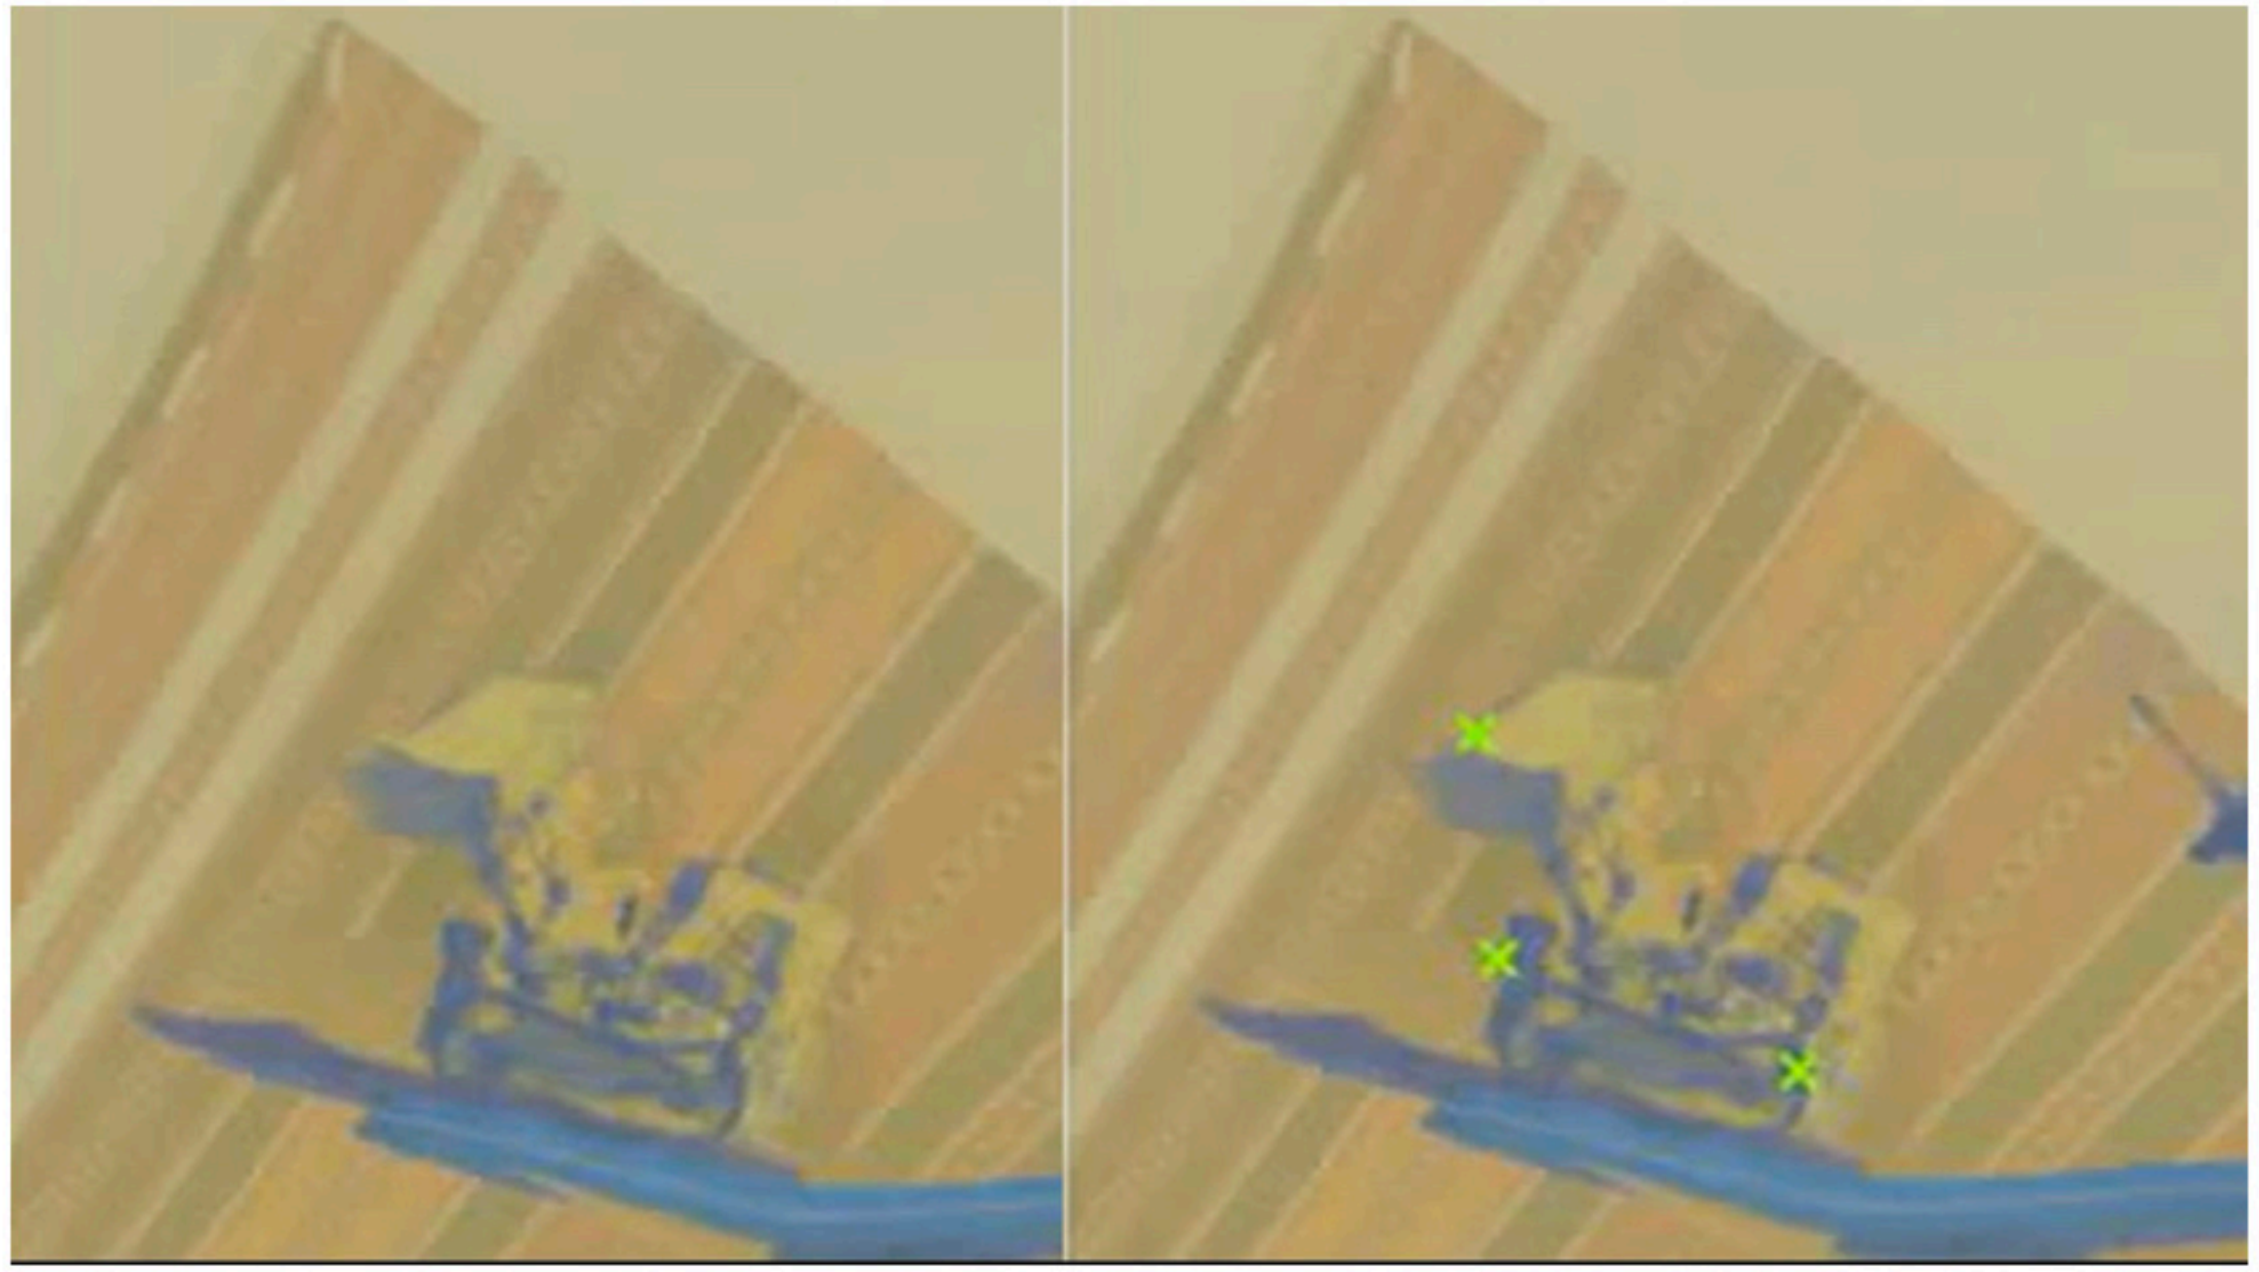

Lily Goli, Daniel Rebain, Sara Sabour, Animesh Garg, Andrea Tagliasacchi “nerf2nerf:Pairwise Registration of Neural Radiance Fields” ICRA 2023

Synthetic Dataset

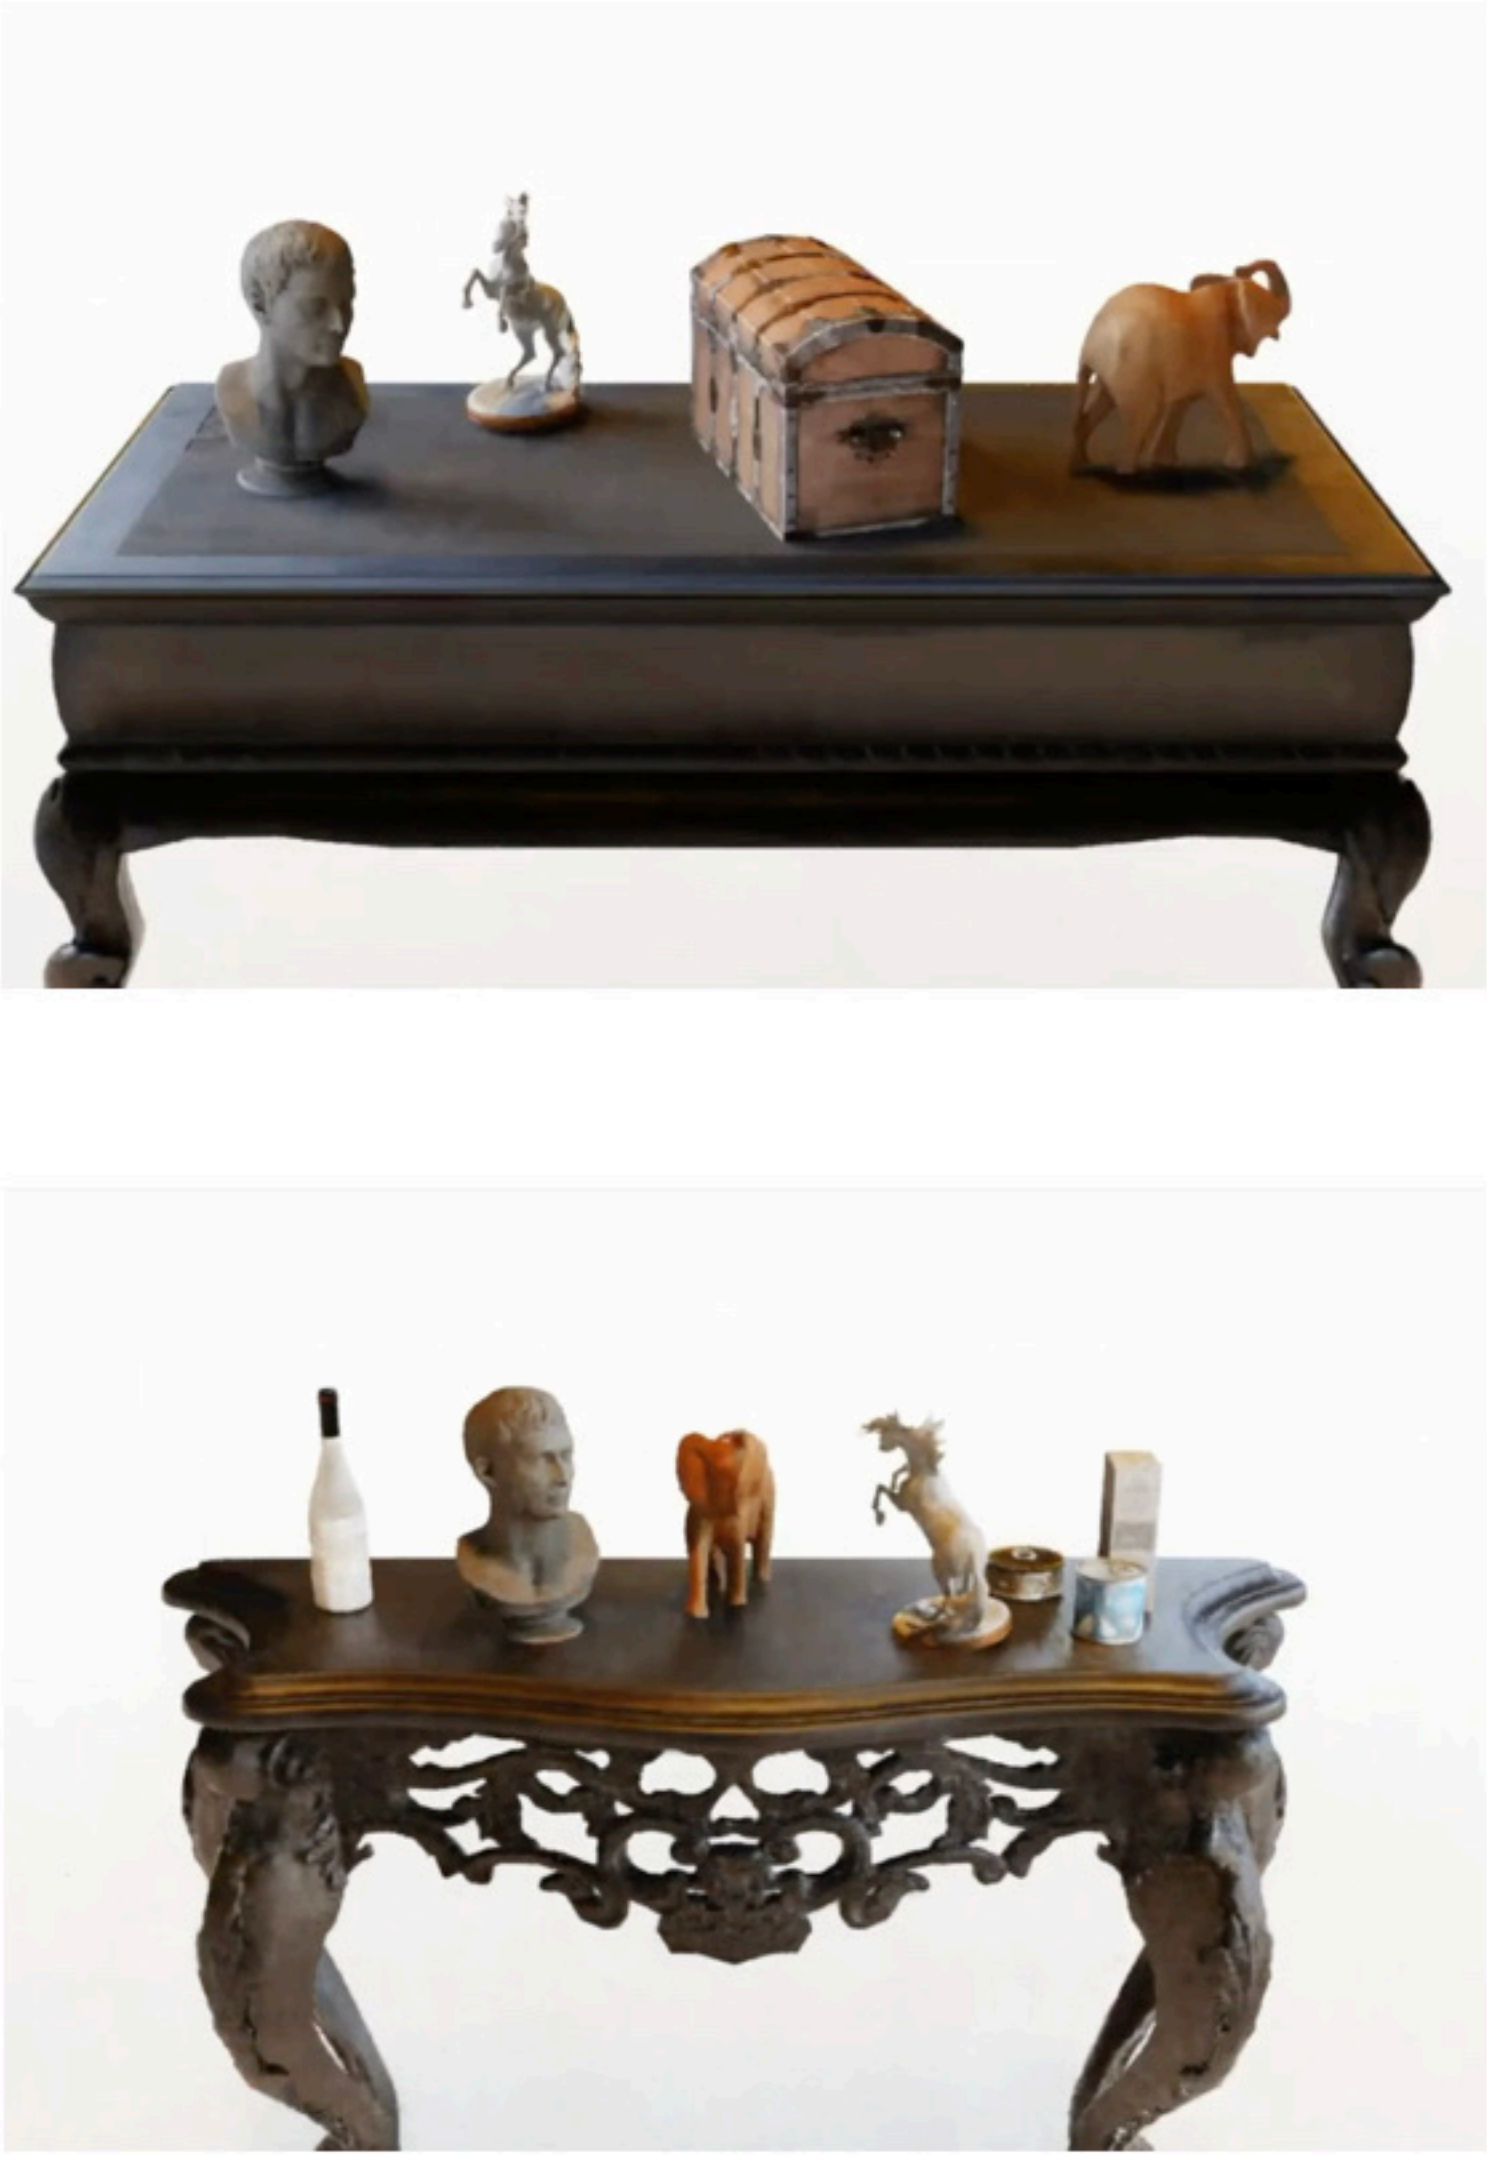

Scene 1

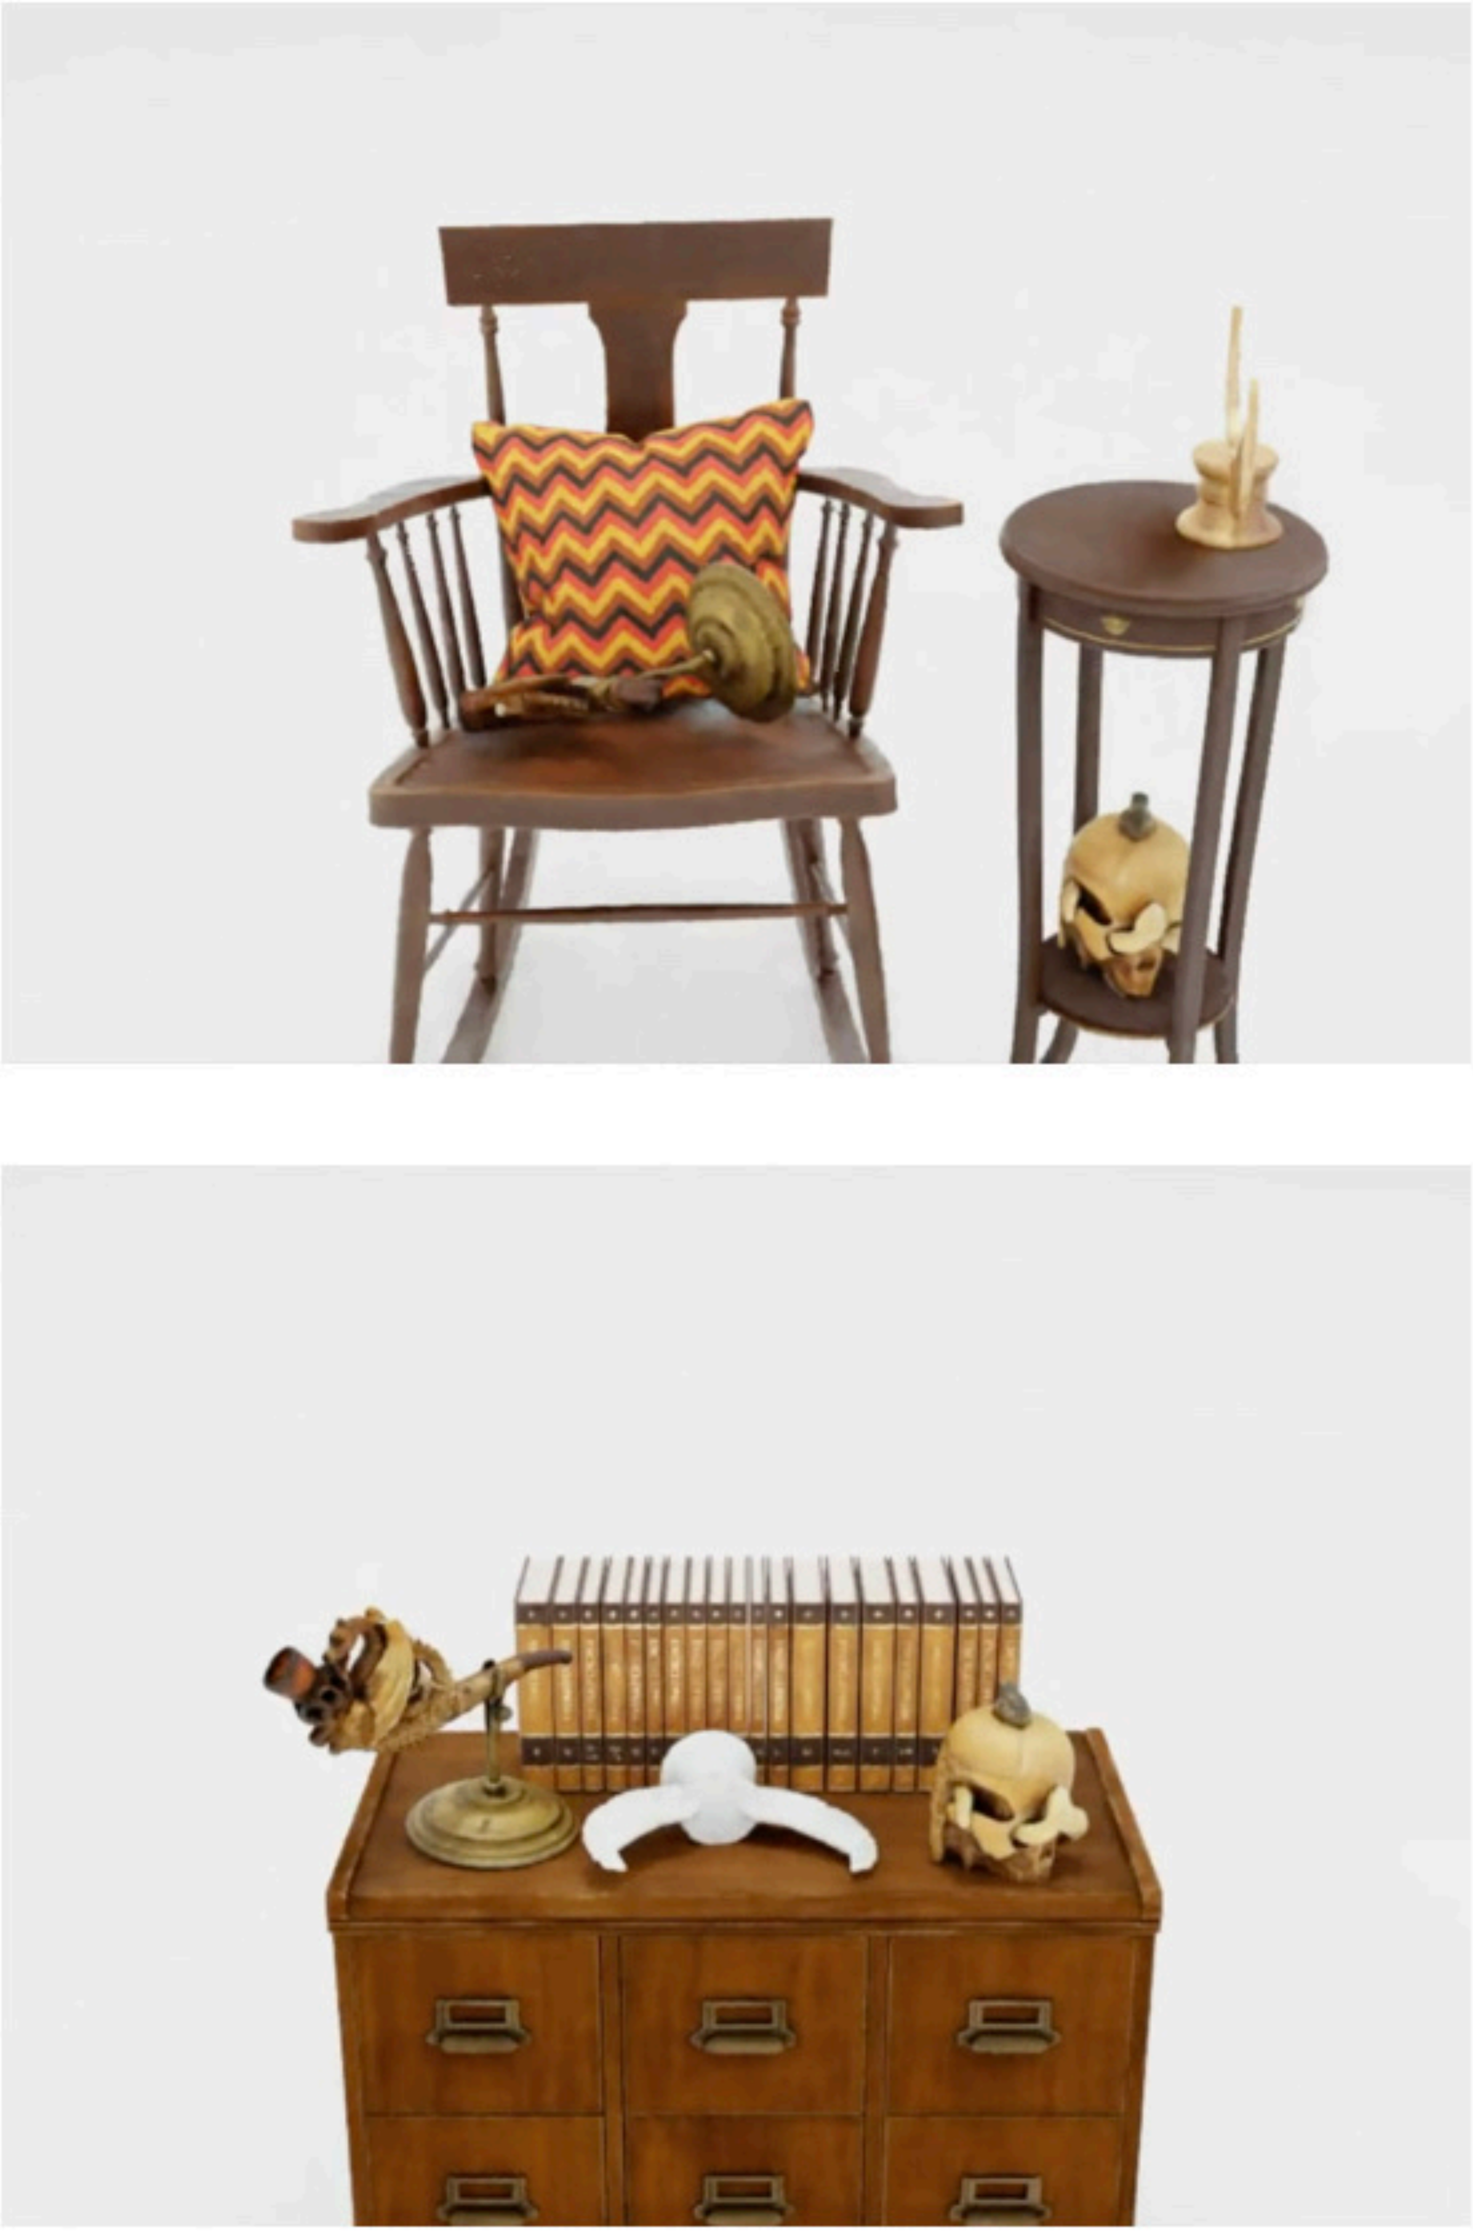

Scene 2

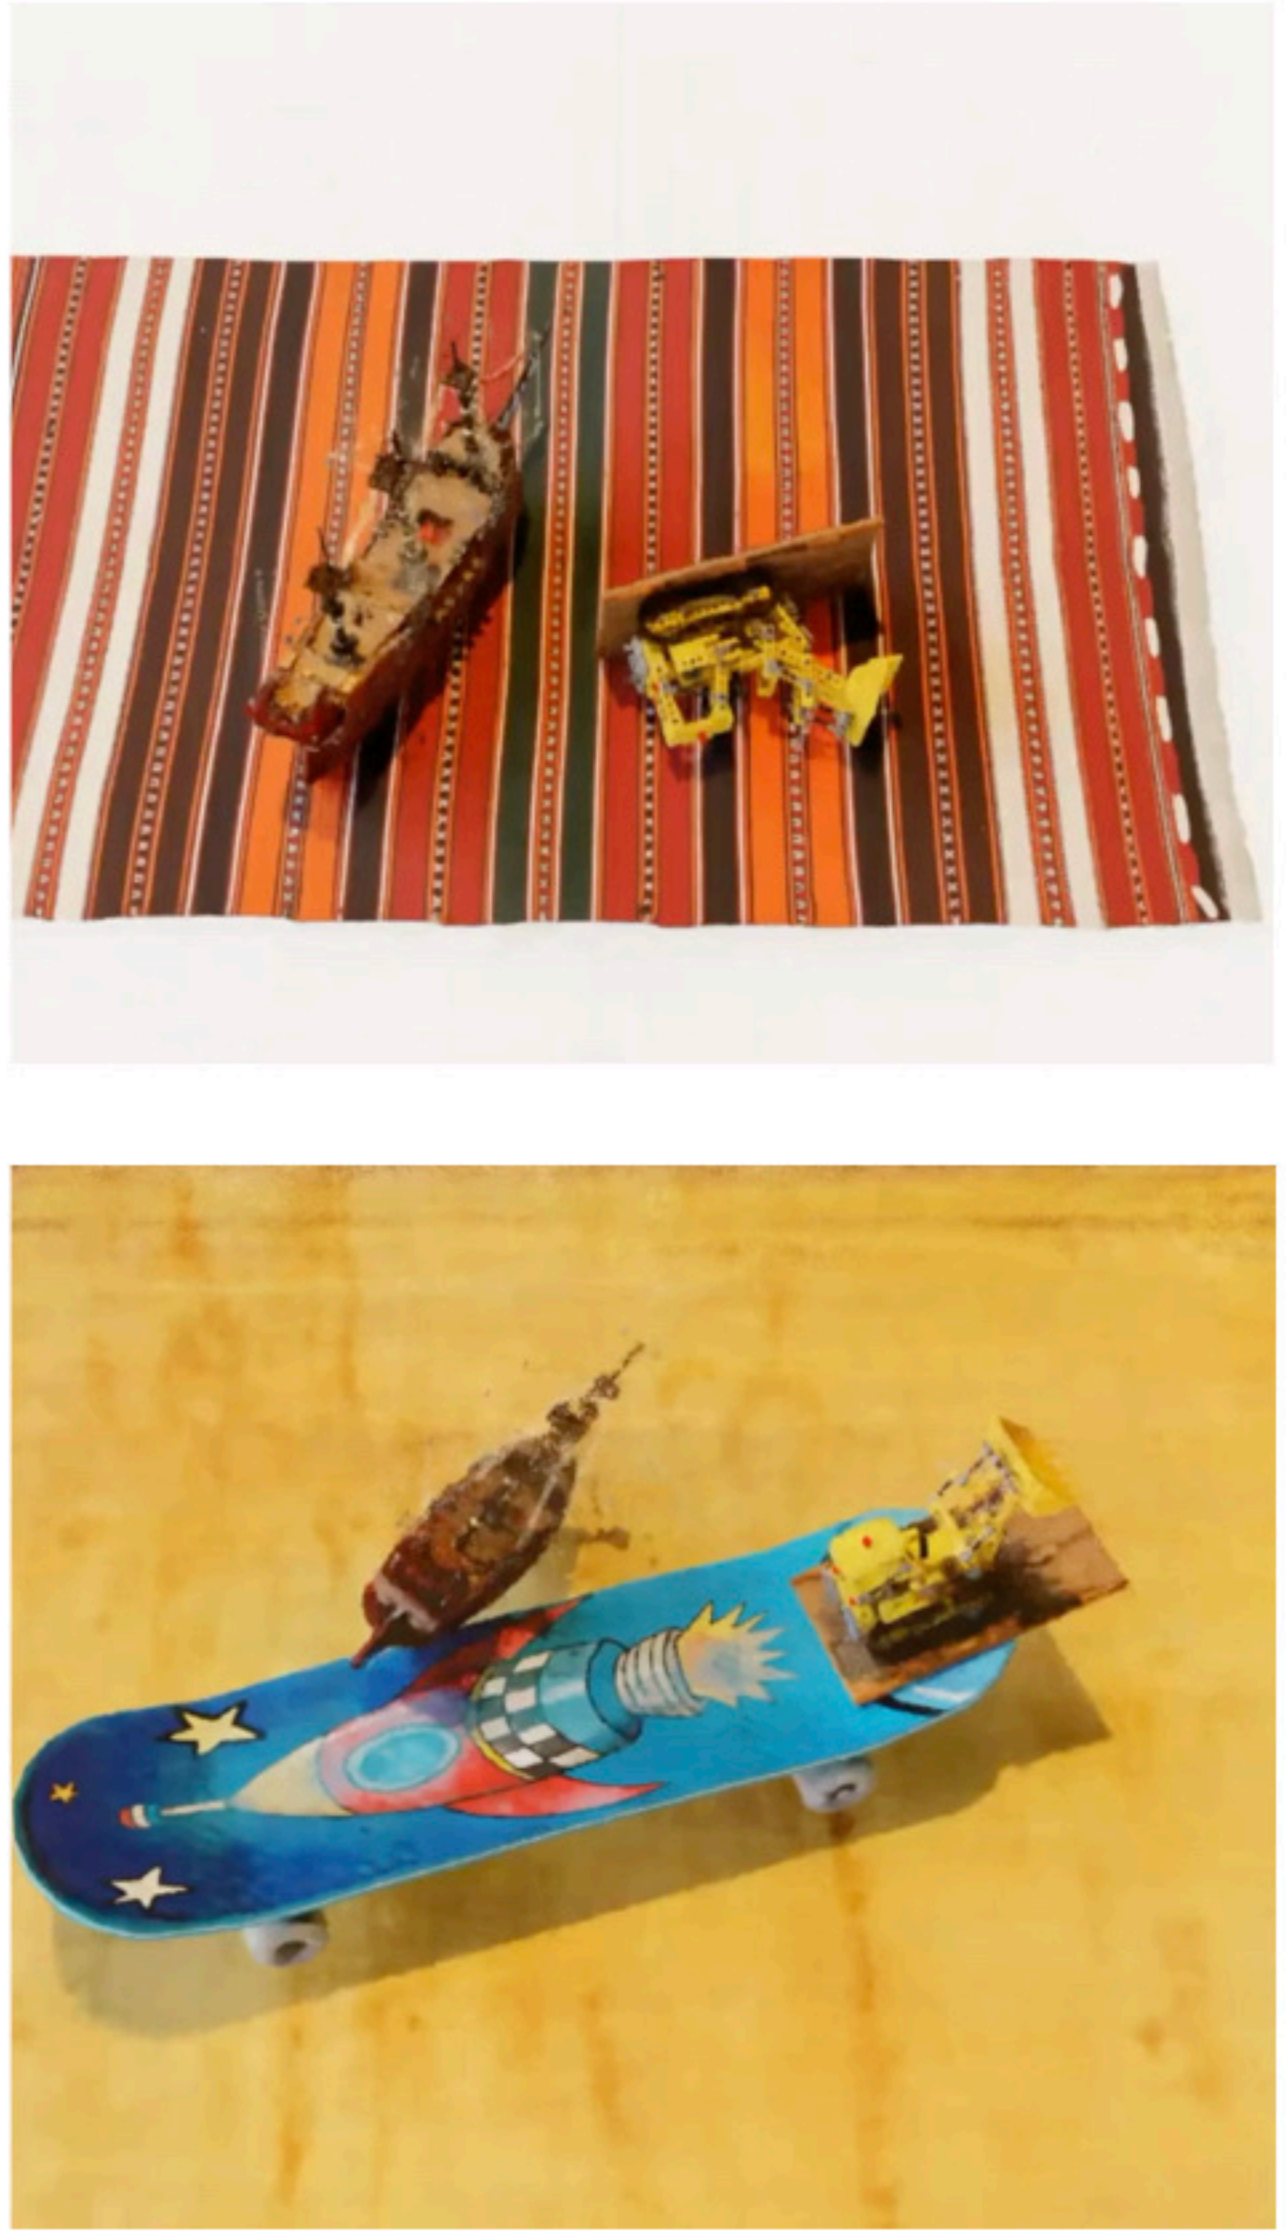

Scene 3

Lily Goli, Daniel Rebain, Sara Sabour, Animesh Garg, Andrea Tagliasacchi “nerf2nerf:Pairwise Registration of Neural Radiance Fields” ICRA 2023

Real Dataset

Outdoor

Indoor

Lily Goli, Daniel Rebain, Sara Sabour, Animesh Garg, Andrea Tagliasacchi "nerf2nerf:Pairwise Registration of Neural Radiance Fields" ICRA 2023

Results (Synthetic)

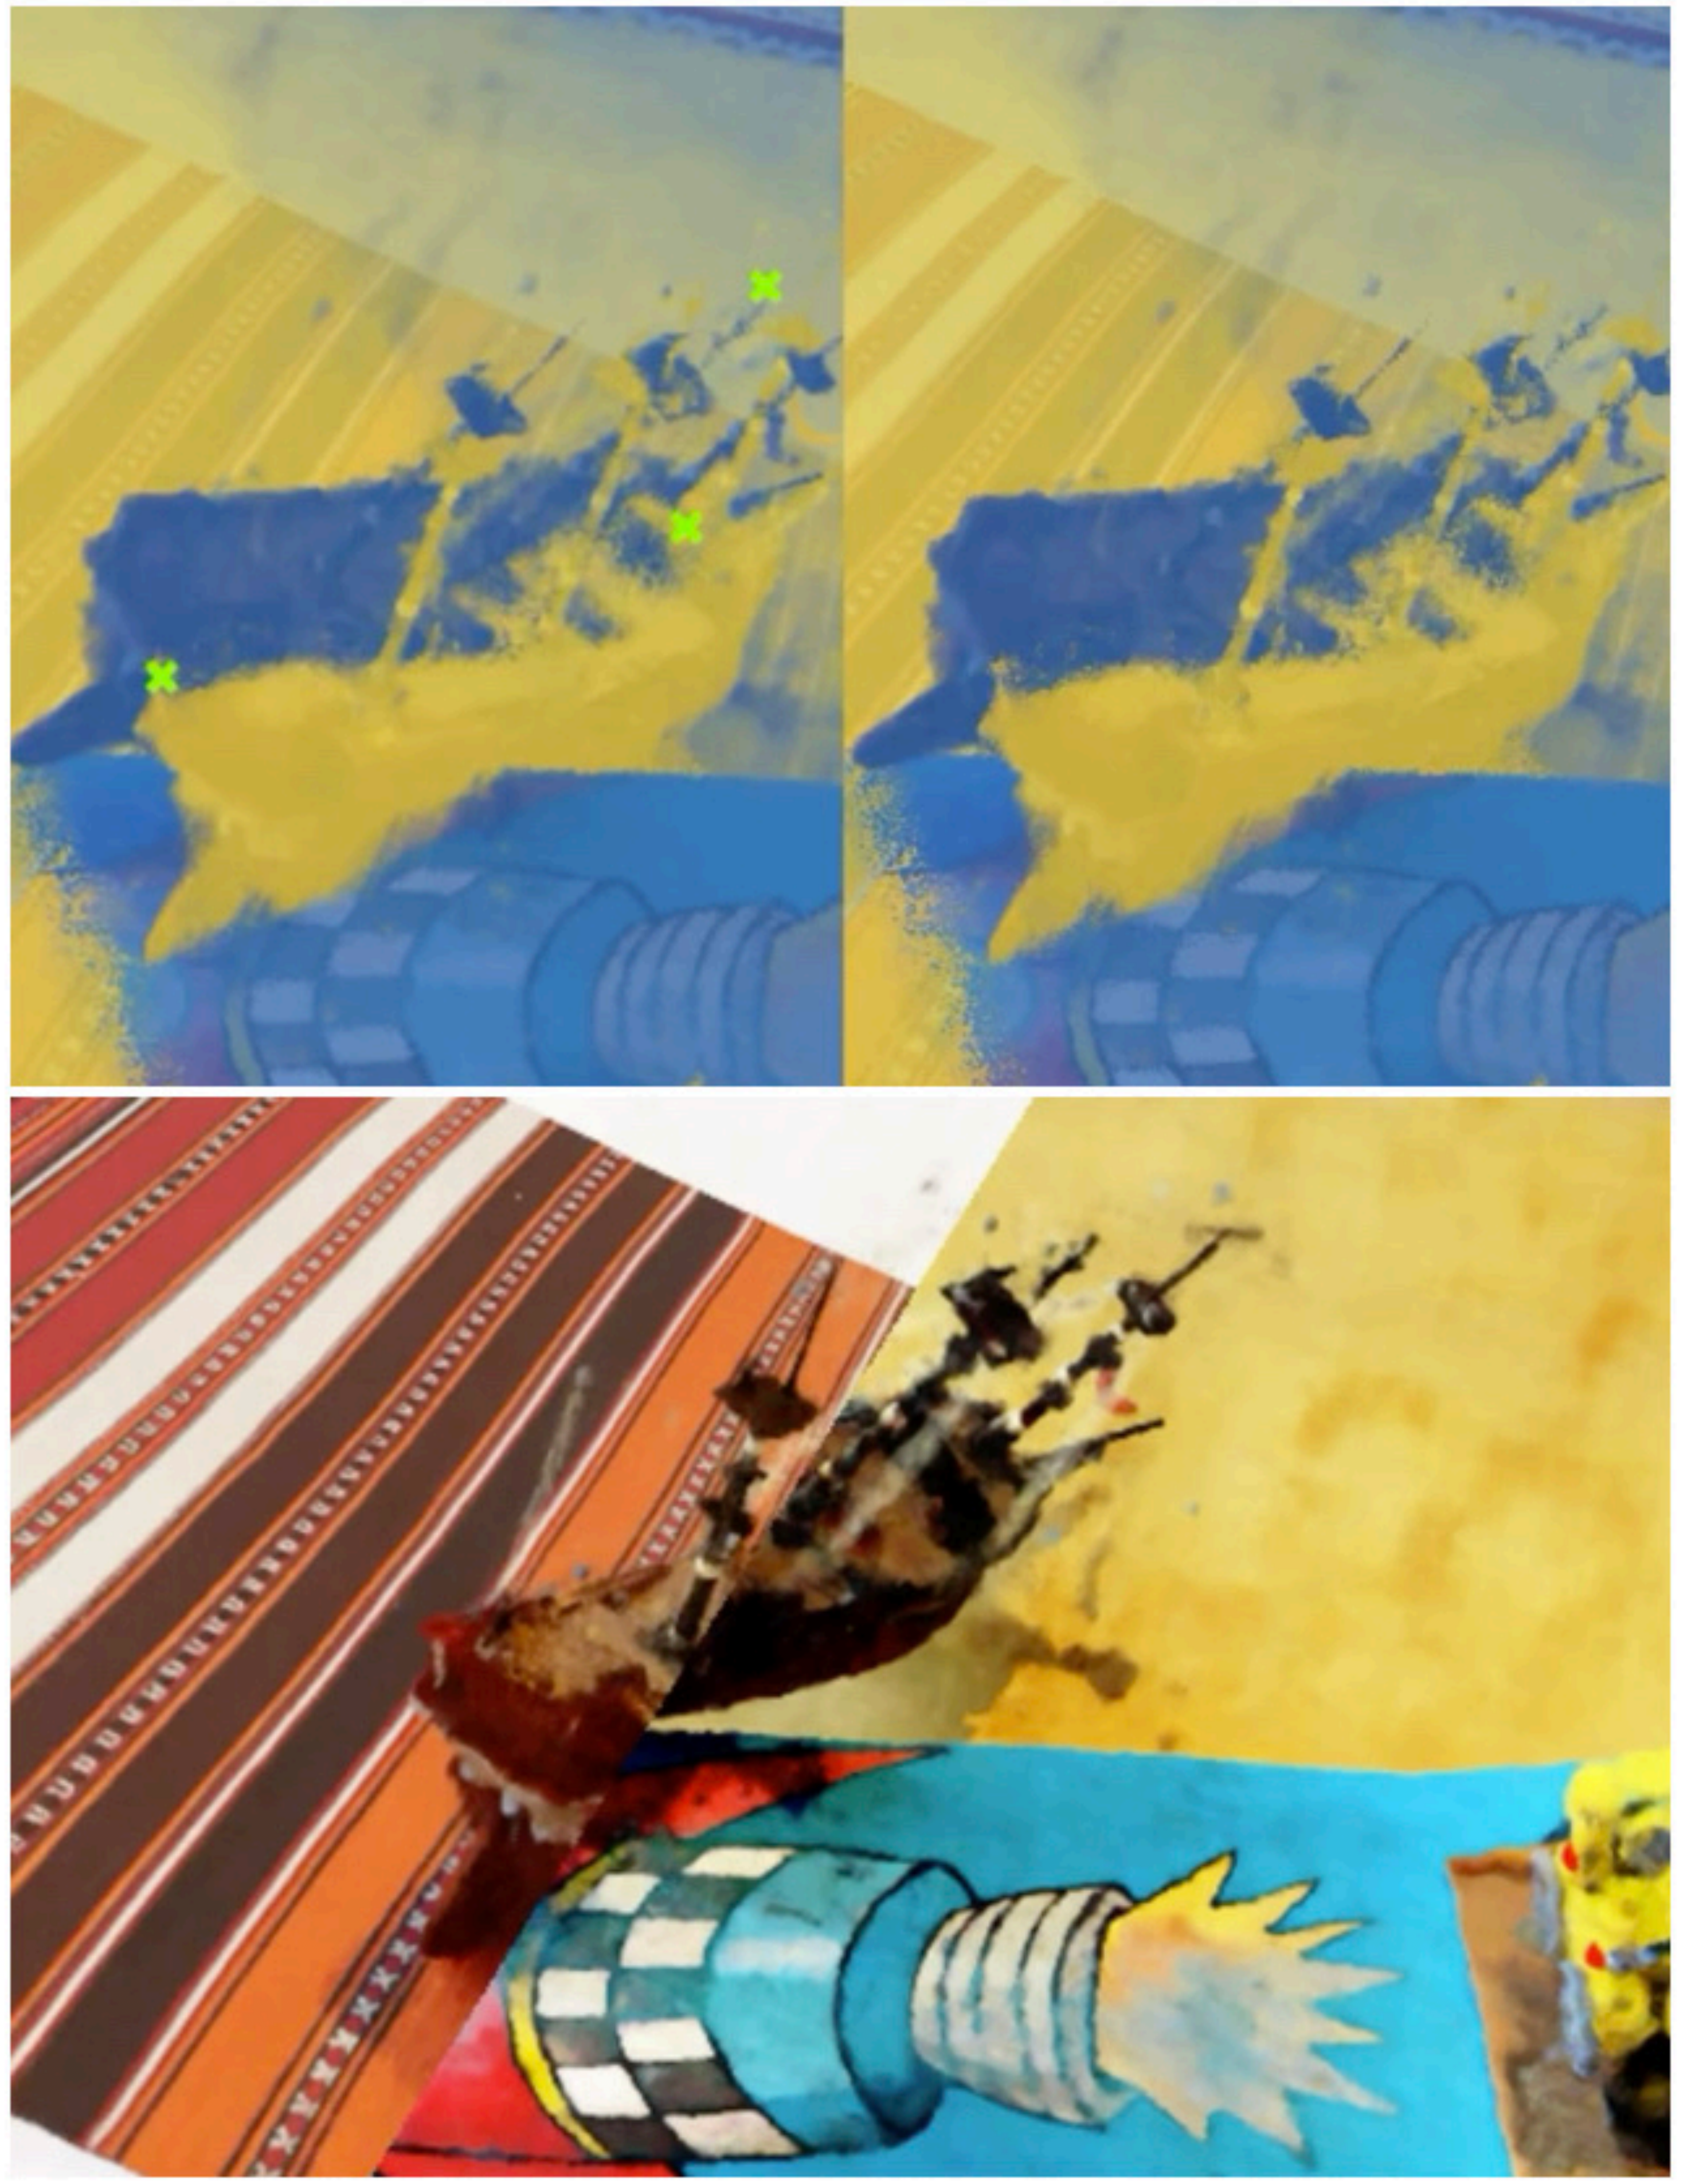

“Ship” $T = 0$

“Pip” $T = 0$

Lily Goli, Daniel Rebain, Sara Sabour, Animesh Garg, Andrea Tagliasacchi “nerf2nerf:Pairwise Registration of Neural Radiance Fields” ICRA 2023

Results (Synthetic)

Object	$10^2 \cdot \Delta t \downarrow$			$\Delta R \downarrow$			$10^2 \cdot$ **3D-ADD** $\downarrow$		
(Pair No.)	KP	FGR	Ours	KP	FGR	Ours	KP	FGR	Ours
bust ①	7.78	7.94	**0.73**	9.93	10.08	**1.39**	7.16	7.02	**0.77**
elephant ①	14.01	12.96	**0.73**	17.16	15.59	**1.00**	13.97	12.87	**0.76**
horse ①	7.88	2.10	**0.99**	14.56	7.32	**1.77**	6.67	2.09	**0.88**
pip ②	4.01	18.93	**0.25**	6.40	30.01	**1.18**	4.61	20.36	**0.46**
jar ②	10.40	7.17	**0.18**	20.46	17.03	**2.84**	11.31	8.13	**0.97**
pedestal ②	5.62	8.13	**0.69**	15.83	11.69	**2.42**	8.42	9.51	**1.23**
lego ③	14.41	12.97	**2.09**	38.08	86.29	**3.89**	18.75	18.40	**2.97**
ship ③	13.48	4.80	**0.75**	20.52	10.74	**1.35**	15.89	5.36	**0.91**

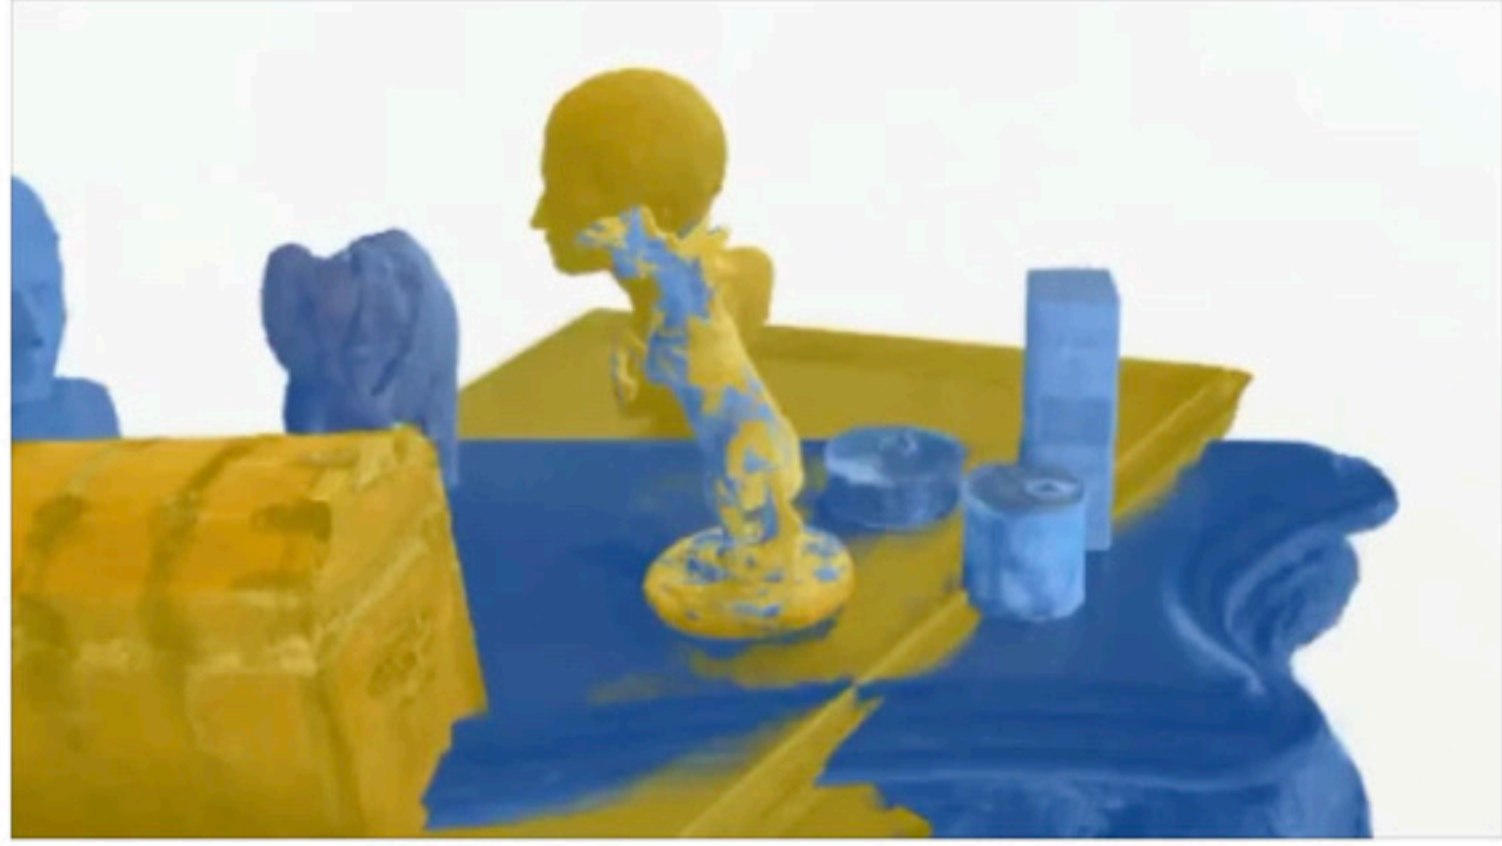

nerf2nerf

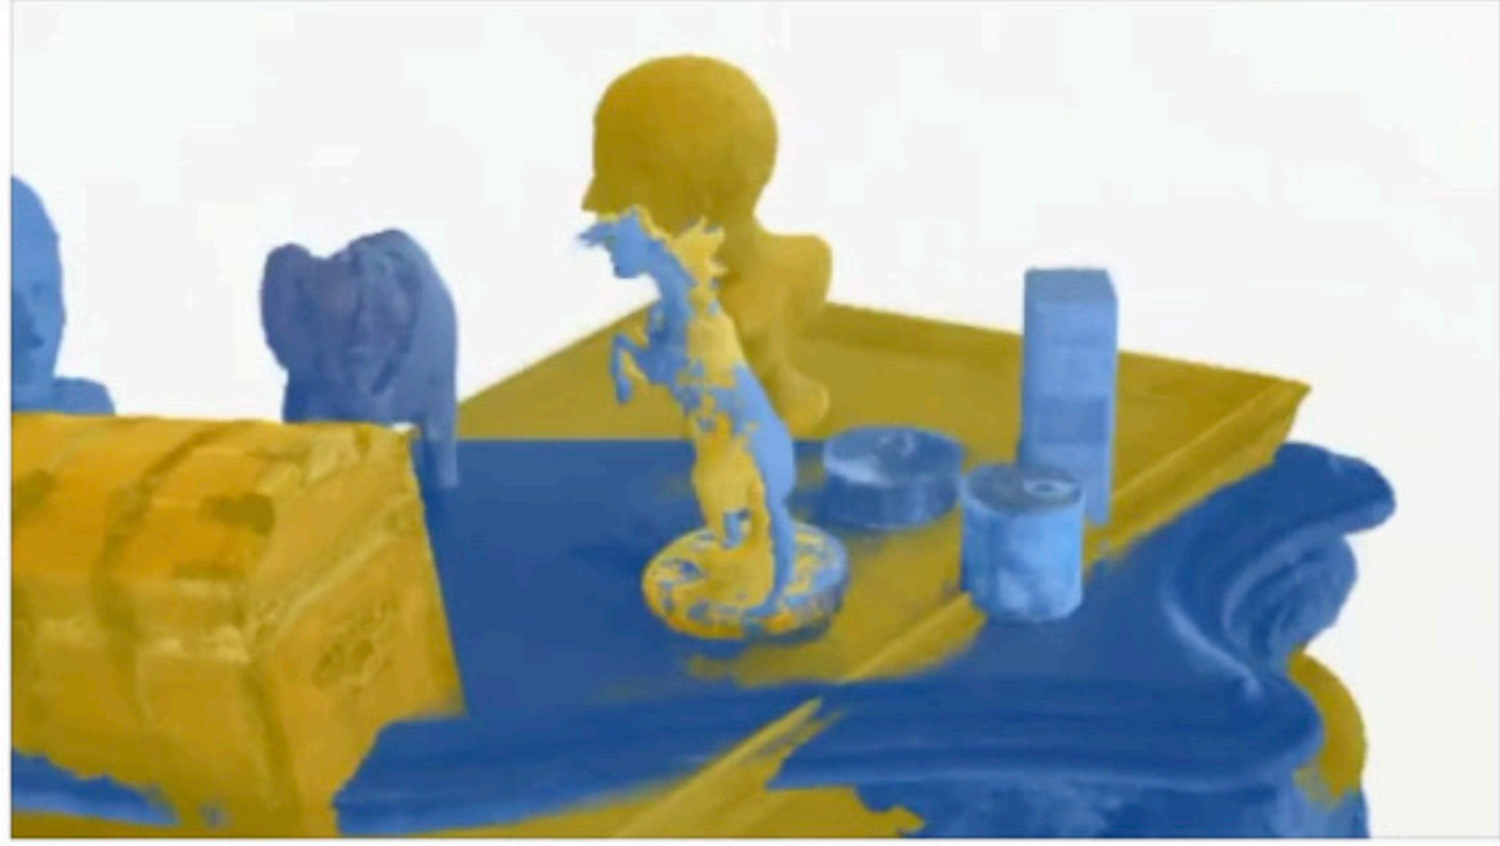

FGR

NeRF Point Cloud

Lily Goli, Daniel Rebain, Sara Sabour, Animesh Garg, Andrea Tagliasacchi “nerf2nerf:Pairwise Registration of Neural Radiance Fields” ICRA 2023

Results (Real)

“Mustard” $T = 0$

“Drill” $T = 0$

Lily Goli, Daniel Rebain, Sara Sabour, Animesh Garg, Andrea Tagliasacchi “nerf2nerf:Pairwise Registration of Neural Radiance Fields” ICRA 2023

An Application

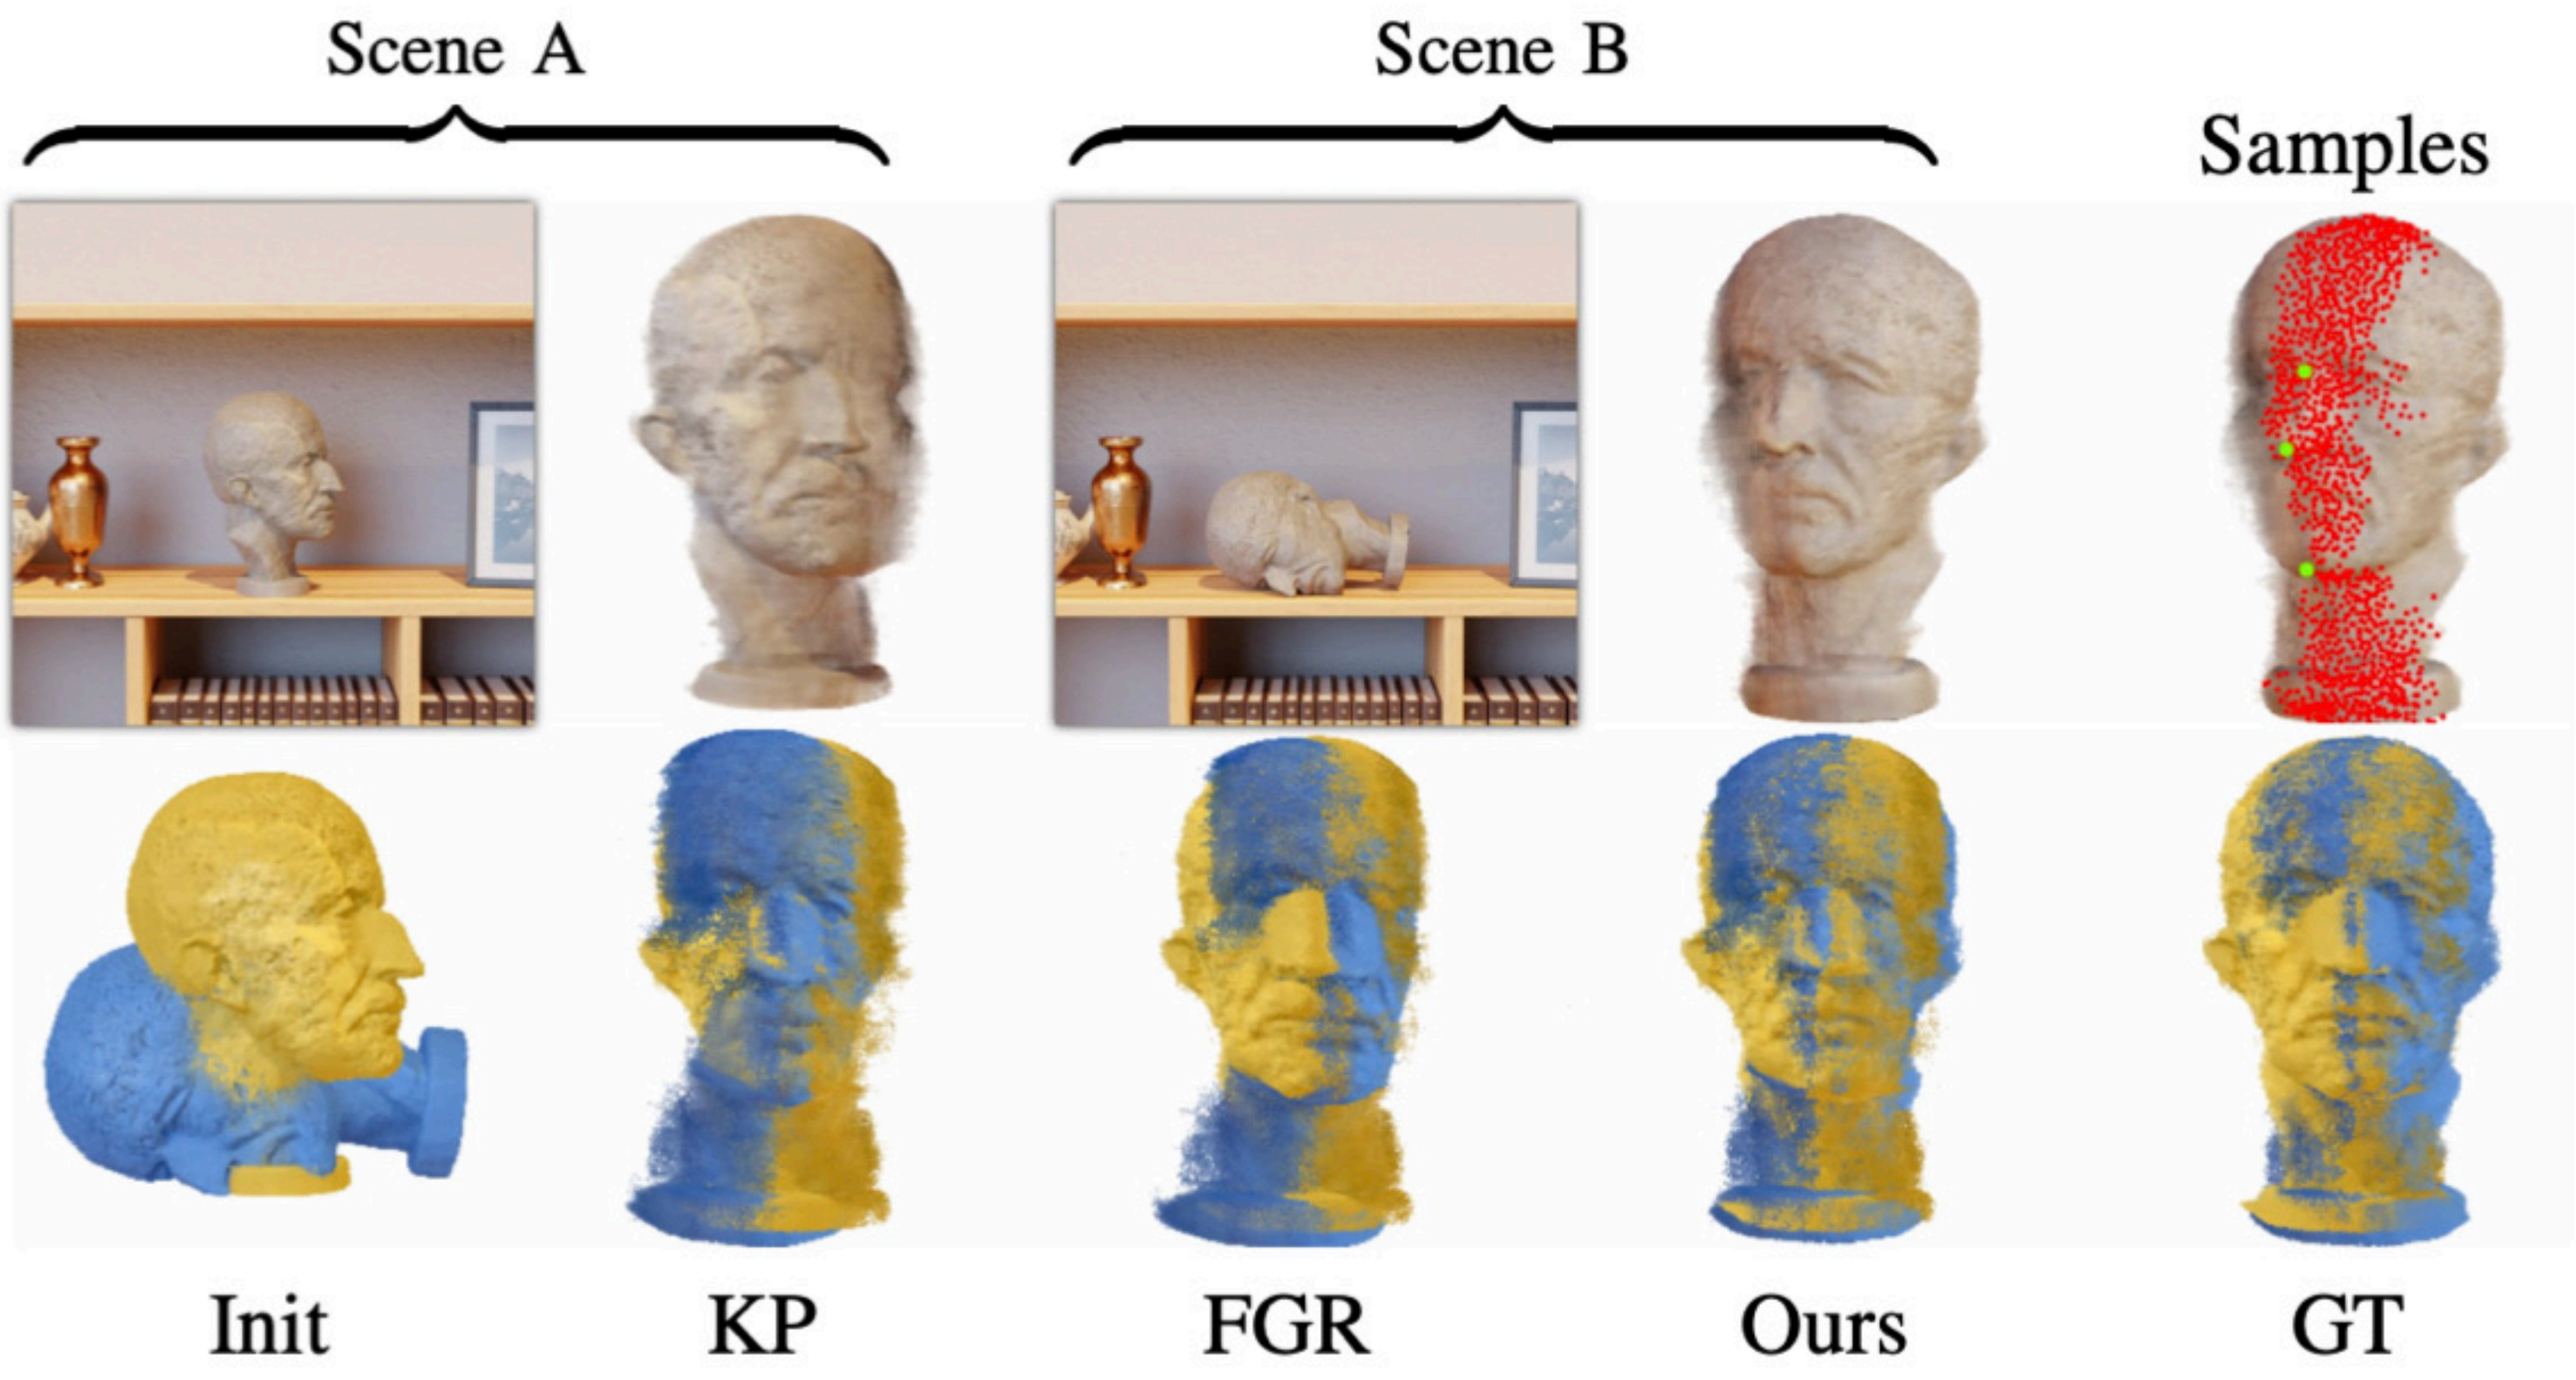

Future work: How to get rid of those *uncertain* parts for a high quality blend of two NeRFs?

(Coming soon!!!)

Goli et al. “nerf2nerf:Pairwise Registration of Neural Radiance Fields” ICRA 2023

Conclusion

- We introduce:
  - a novel geometric representation extracted from NeRF called "Surface Fields"
  - a rigid registration method between two neural radiance fields (NeRFs)
  - a dataset of paired NeRFs of real and synthetic scenes with differently posed multiple objects
- Limitations:
  - Sampling can fail on thin or small objects
  - Registering geometrically symmetric objects with asymmetric texture

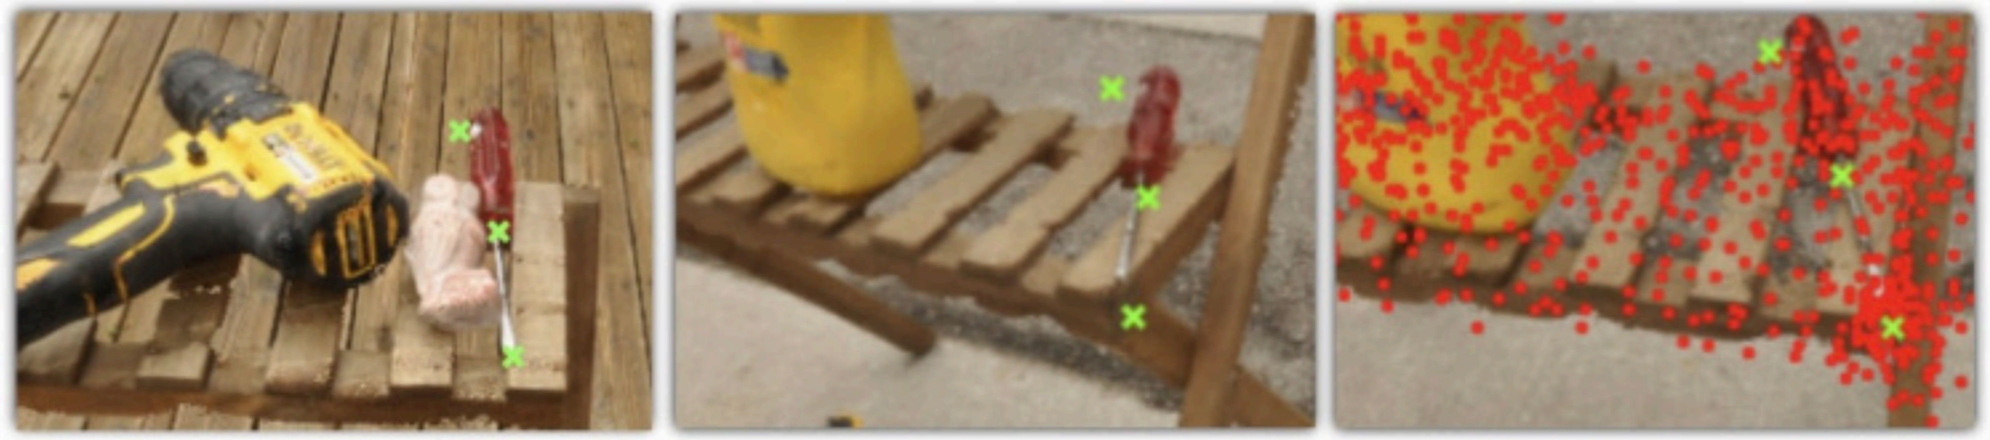

Scene A Scene B Samples

Lily Goli, Daniel Rebain, Sara Sabour, Animesh Gary, Andrea Tagliassachi "nerf2nerf:Pairwise Registration of Neural Radiance Fields" ICRA 2023

nerf2nerf: Pairwise Registration of Neural Radiance Fields

Lily Goli[1,2], Daniel Rebain[4], Sara Sabour[1,2,6], Animesh Garg[1,2,5], Andrea Tagliasacchi[1,3,6]

[1]University of Toronto, [2]Vector Institute, [3]SFU, [4]UBC, [5]NVIDIA, [6]Google Research

https://nerf2nerf.github.io