ICCV HONOLULU HAWAII
OCT 19-23, 2025

UnMix-NeRF: Spectral Unmixing Meets Neural Radiance Fields


Fabian Perez


Sara Rojas


Carlos Hinojosa


Hoover Rueda-Chacón


Bernard Ghanem

Universidad Industrial de Santander

King Abdullah University of Science and Technology

NeRF – Neural Radiance Fields

Synthesizes novel views of 3D scenes from 2D images by representing the scene as networks

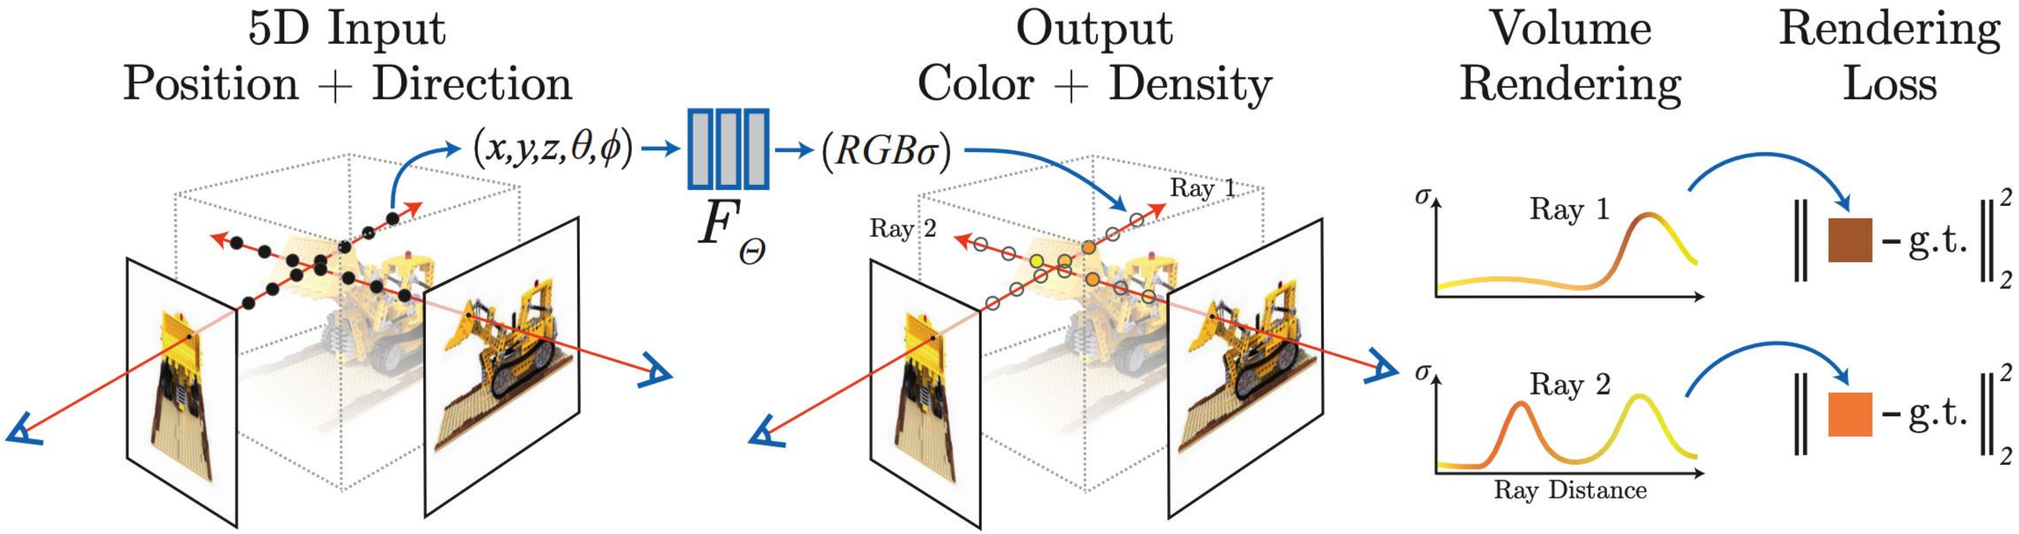

[1] Mildenhall, Ben, et al. "Nerf: Representing scenes as neural radiance fields for view synthesis." *Communications of the ACM* 65.1 (2021): 99-106.

NeRF – Neural Radiance Fields

Synthesizes novel views of 3D scenes from 2D images by representing the scene as networks



- Limited to RGB domain
- Traditional NeRFs capture geometry and appearance but **lack material information**
- Needs external features/models for any kind of segmentation

Limitations on RGB

Metamerism

Two different materials match under one light but mismatch under another

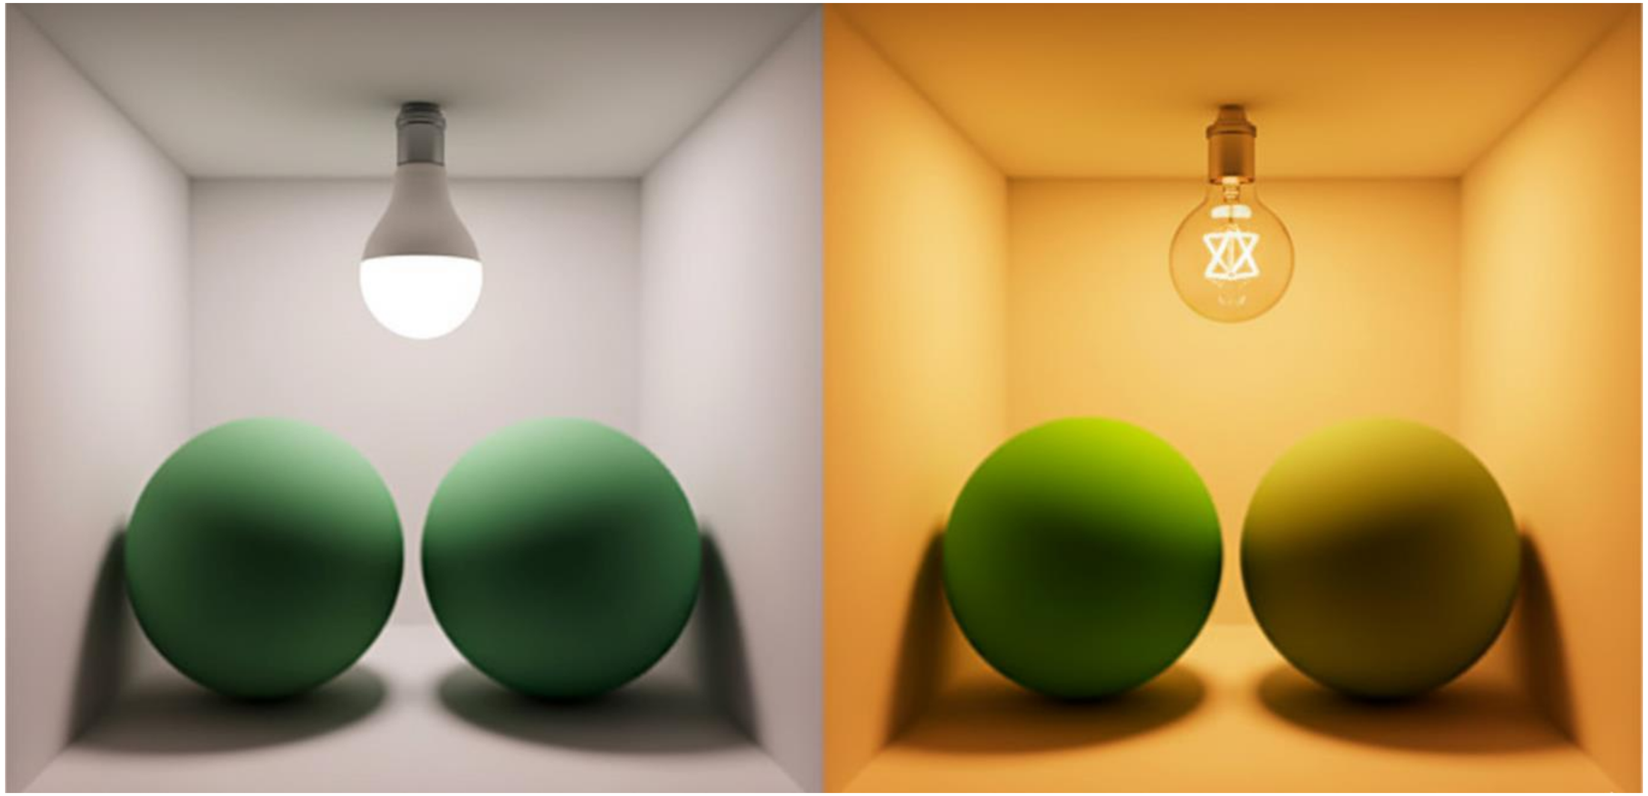

Limitations on RGB

Why material understanding matters?

Metamerism

Two different materials match in one light, mismatch in another

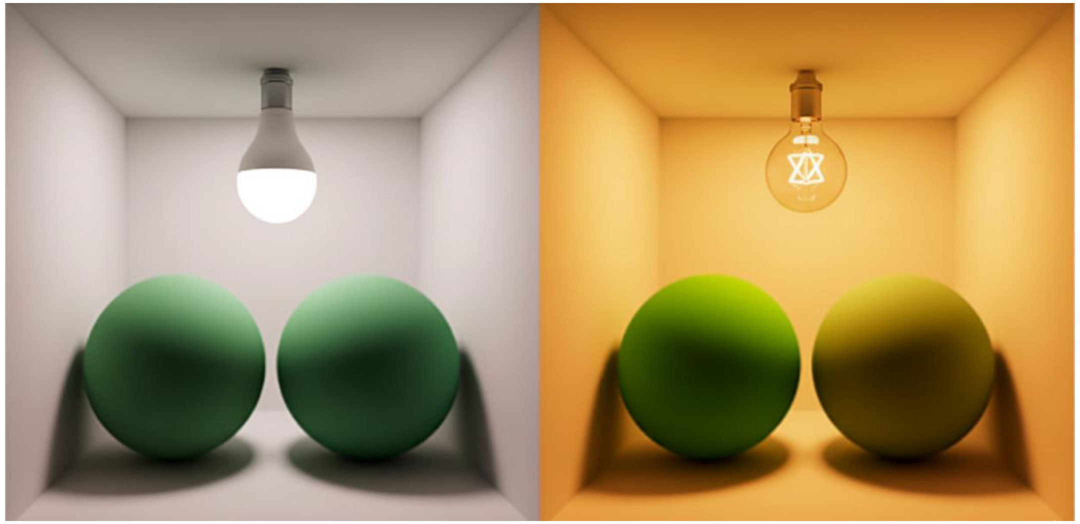

Robotics

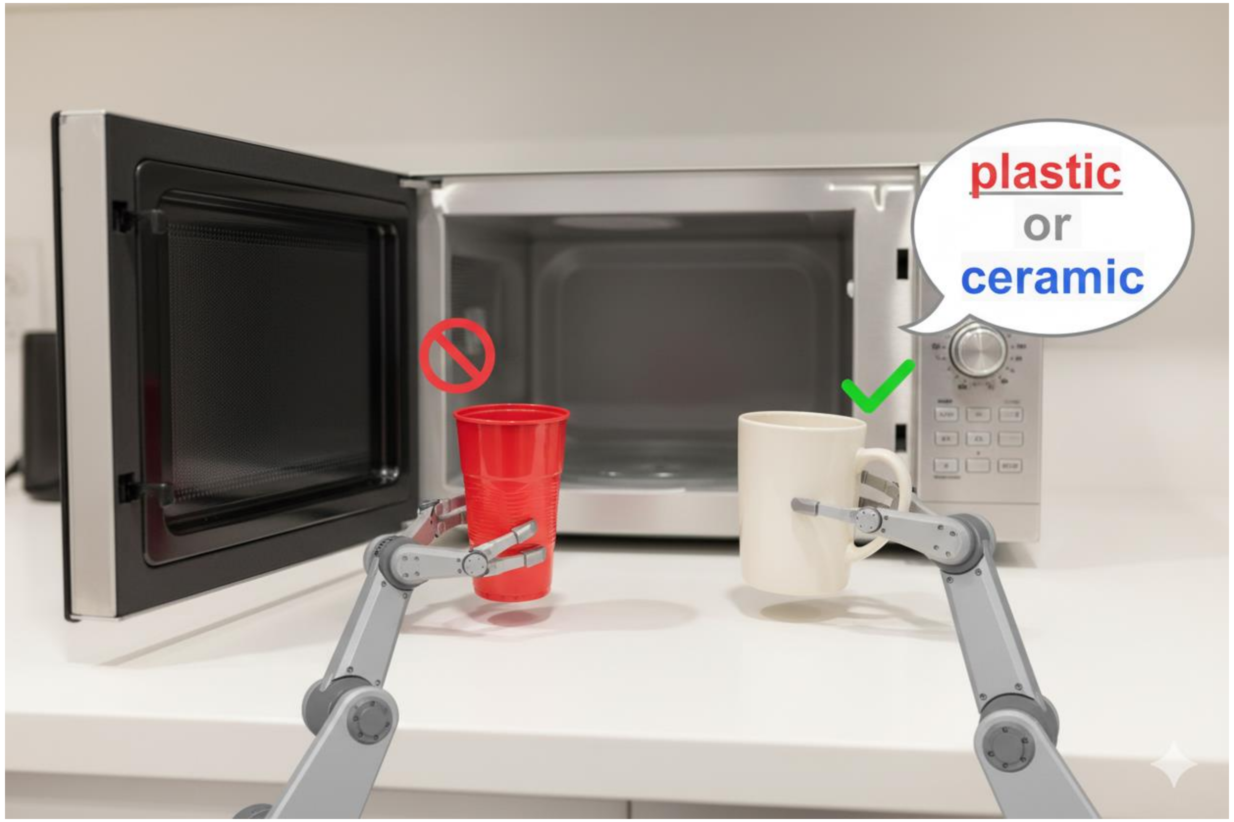

Spectral Imaging

Captures images beyond the visible by sampling light in many bands, so each pixel encodes a rich spectral profile

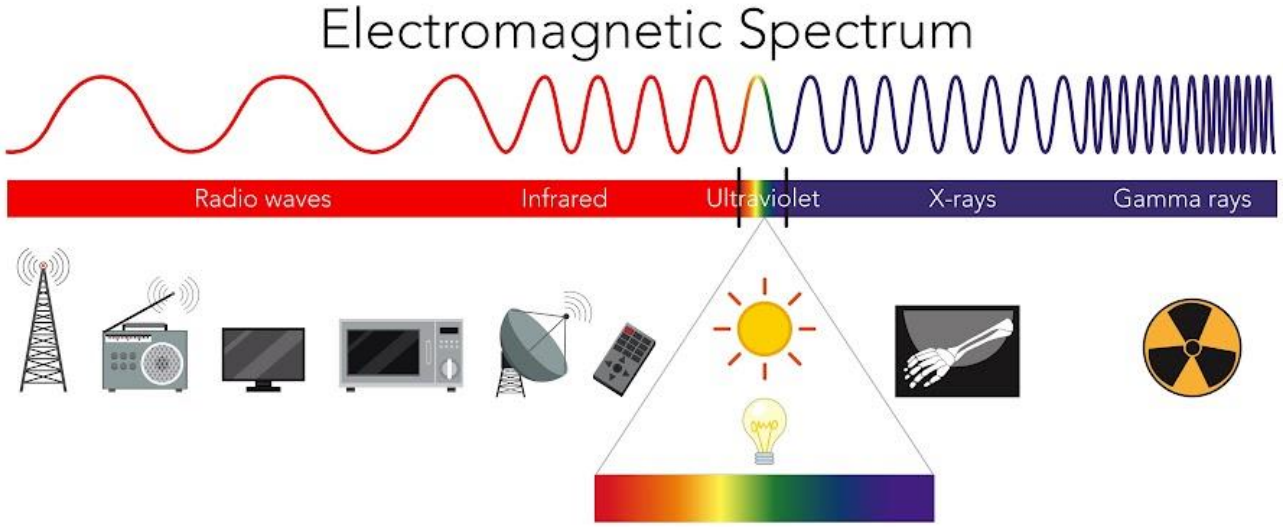

More bands create a spectral cube (x, y, λ) instead of a 3-channel image

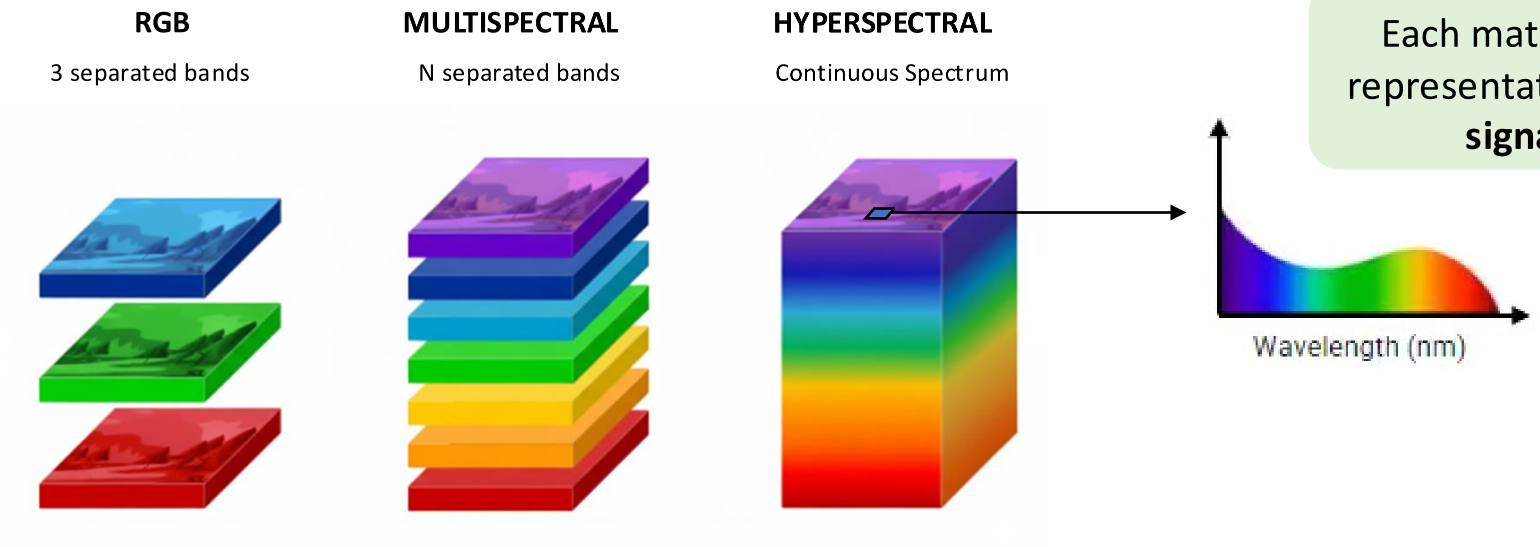

Each material has a representative **spectral signature**

Spectral Imaging: Applications

Material Segmentation

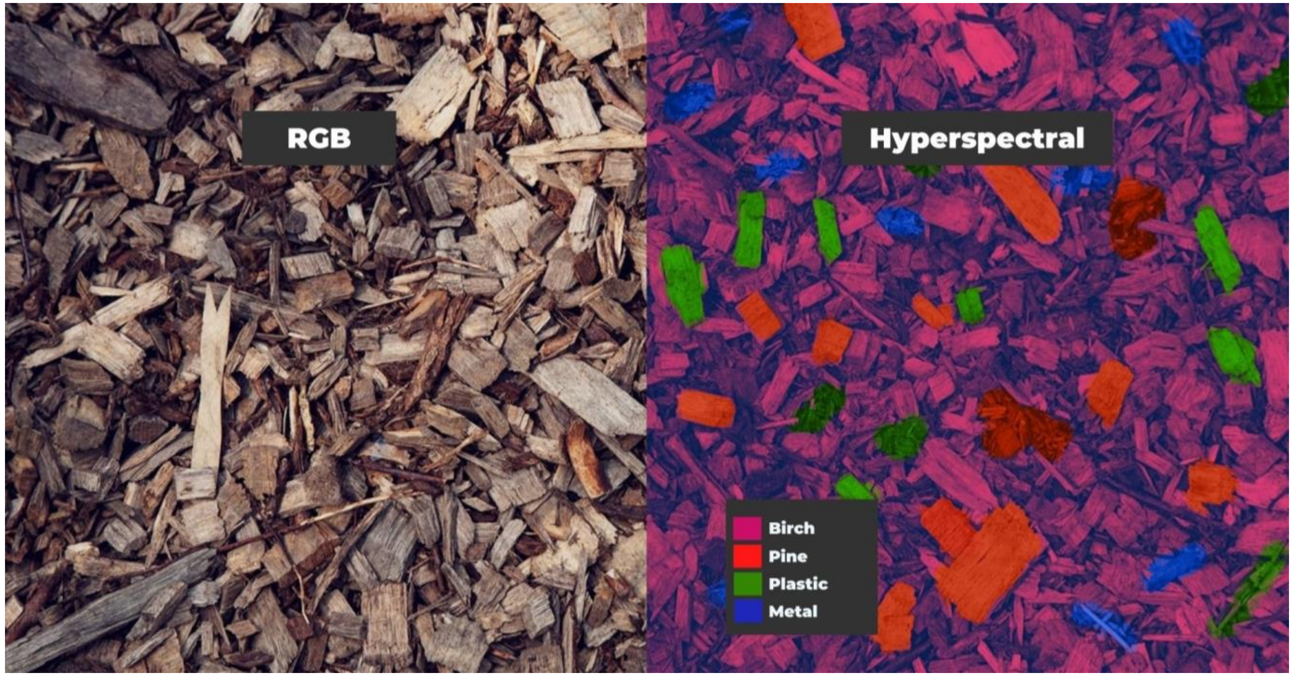

Classification

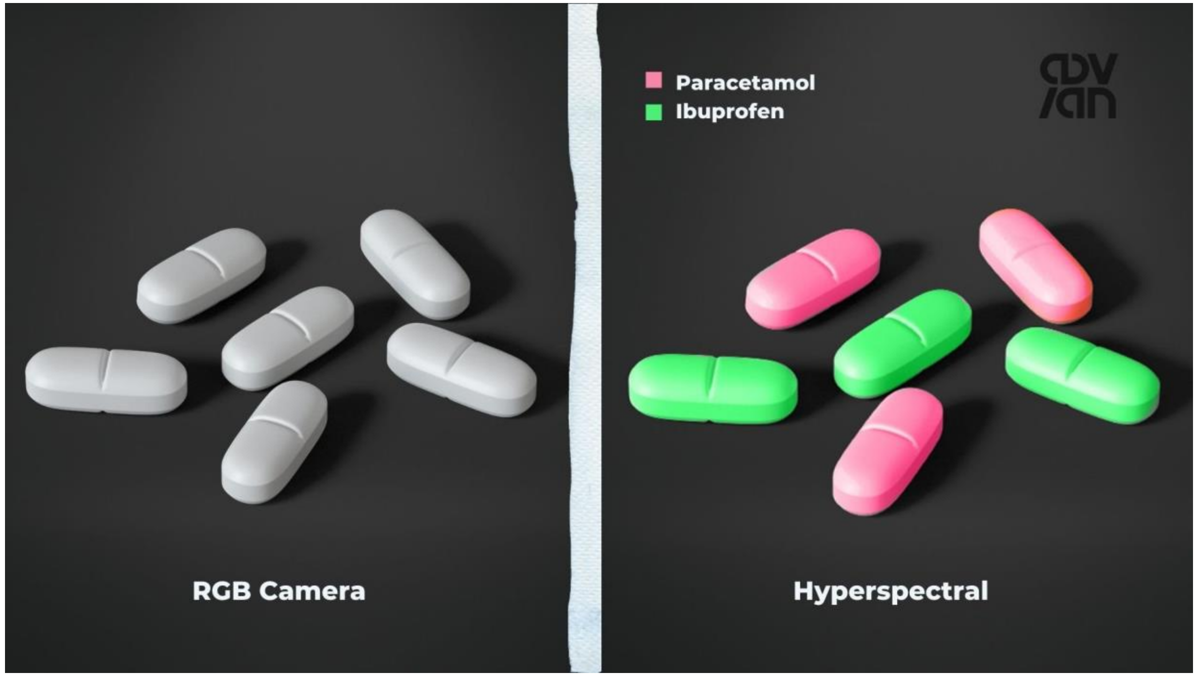

[1] https://www.advian.fi/en/what-is-hyperspectral-imaging

Spectral Unmixing

Decomposes each pixel's spectrum into materials (endmembers) and their abundances

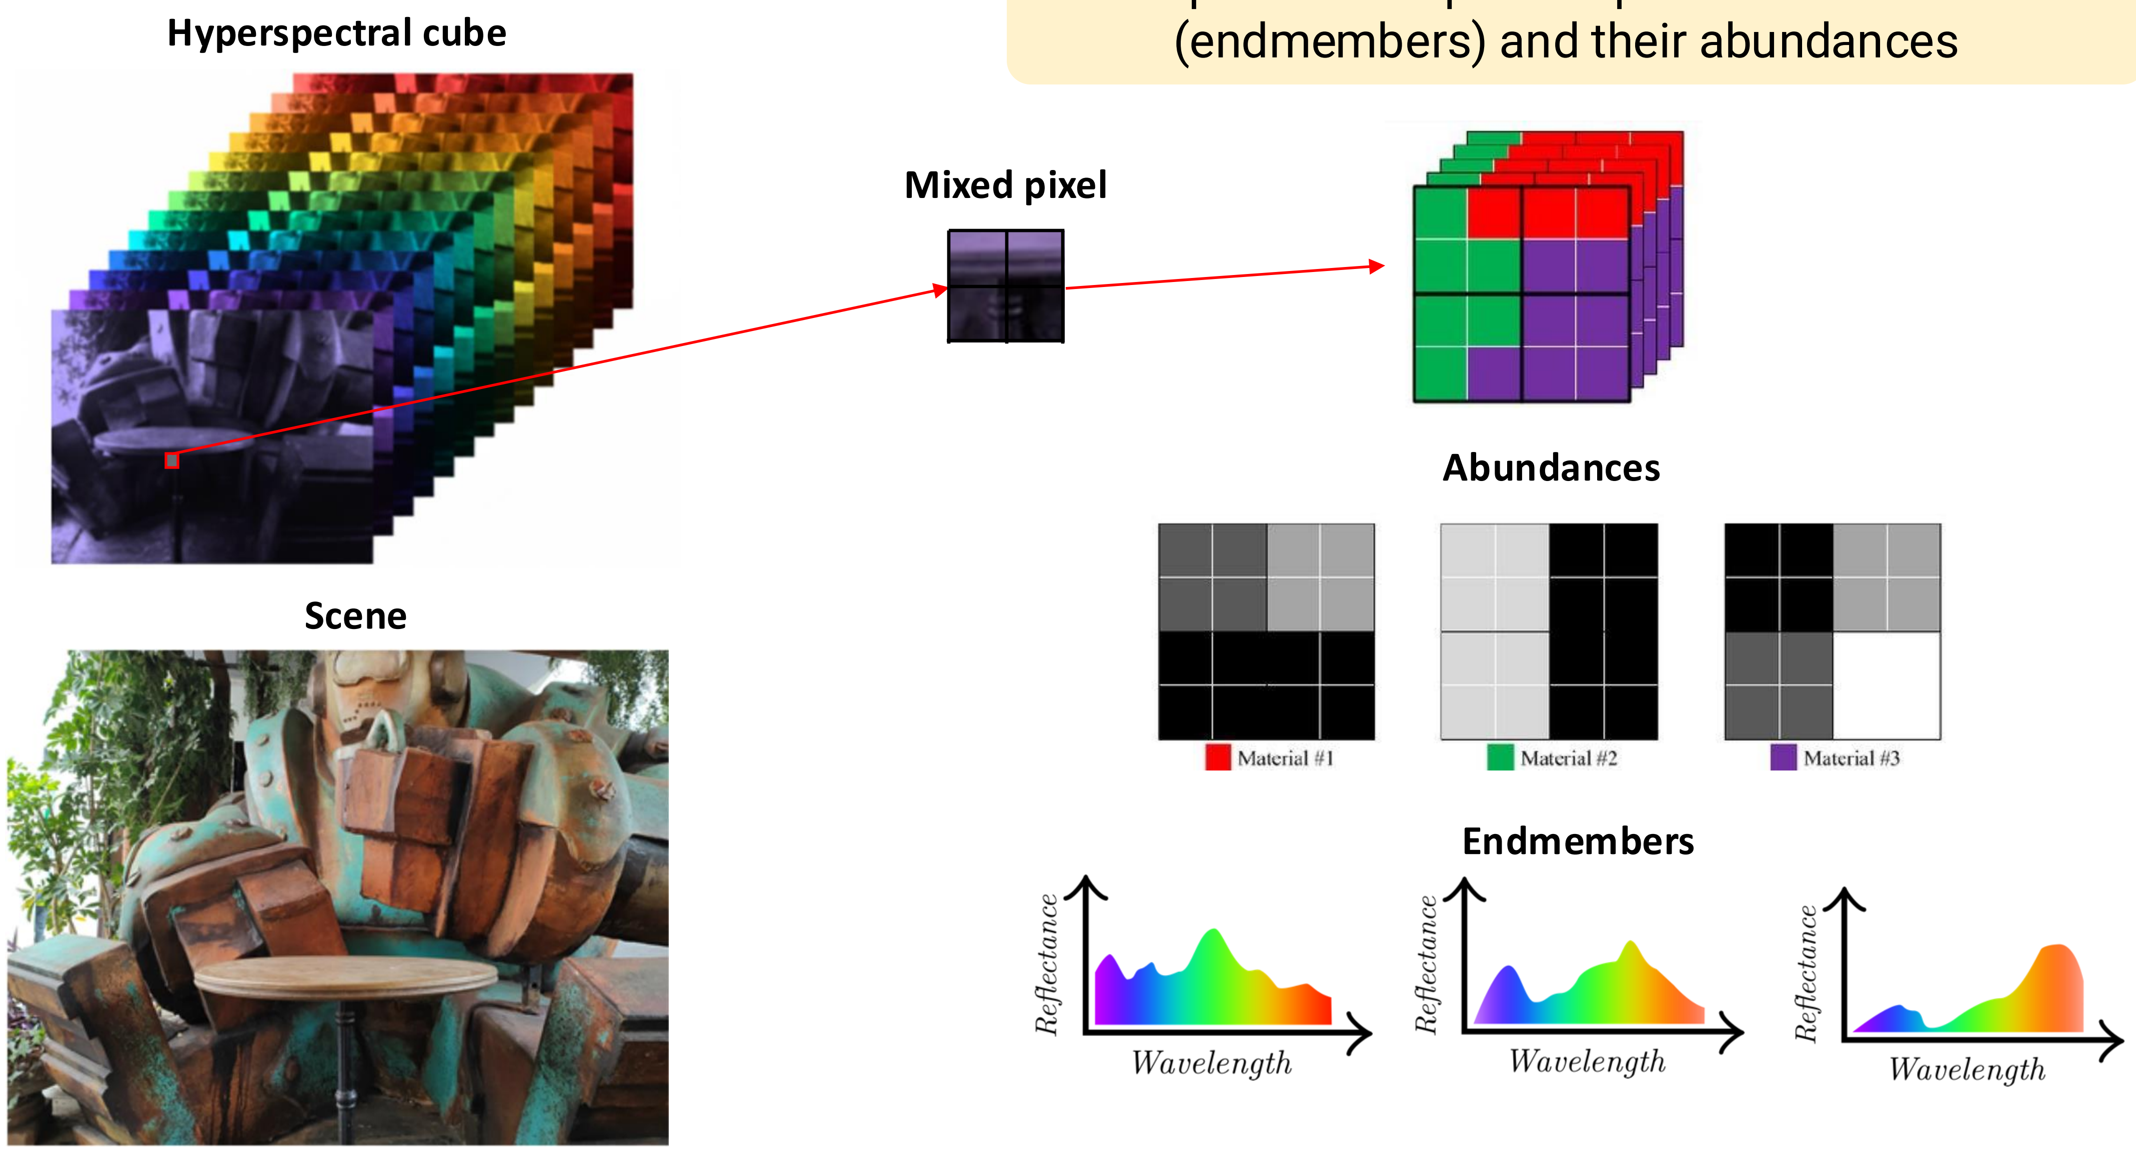

Spectral Unmixing

Decomposes each pixel's spectrum into materials (endmembers) and their abundances

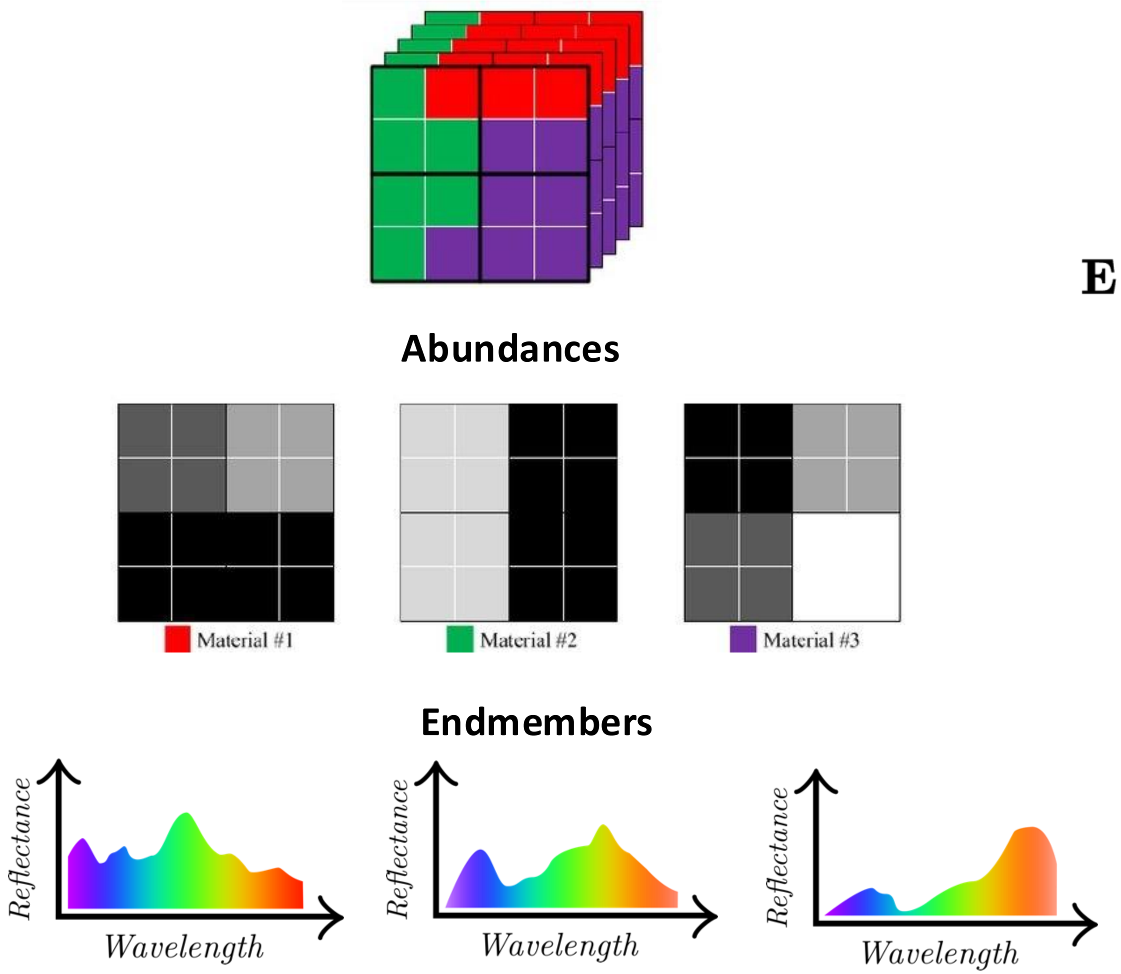

Each spectral pixel is a linear combination of:

$$\mathbf{y}_n = \mathbf{E}\mathbf{a}_n + \boldsymbol{\epsilon}_n$$

Endmembers Matrix

$$\mathbf{E} = [\mathbf{e}_1, \cdots, \mathbf{e}_k, \cdots, \mathbf{e}_K] \in \mathbb{R}^{B \times K}$$

Abundances vector

$$\mathbf{a}_n \in \mathbb{R}^K$$

The following constraints must be satisfied:

Non-negativity

$$\mathbf{a}_n \succeq 0$$

Sum-to-One

$$\mathbf{1}^T \mathbf{a}_n = 1$$

Box constraint

$$\forall\, i, j, \quad 0 \leq e_{ij} \leq 1$$

3D Method with Spectral Imaging

HS-NeRF [1]

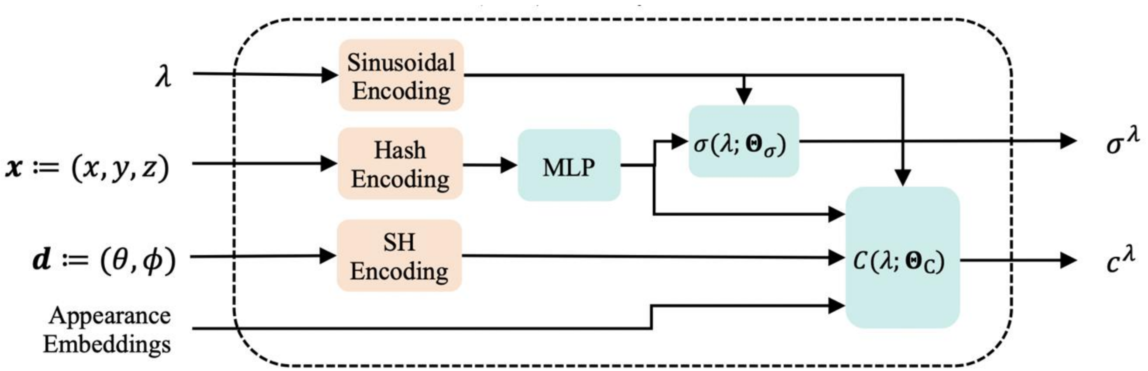

SpectralNeRF [2]

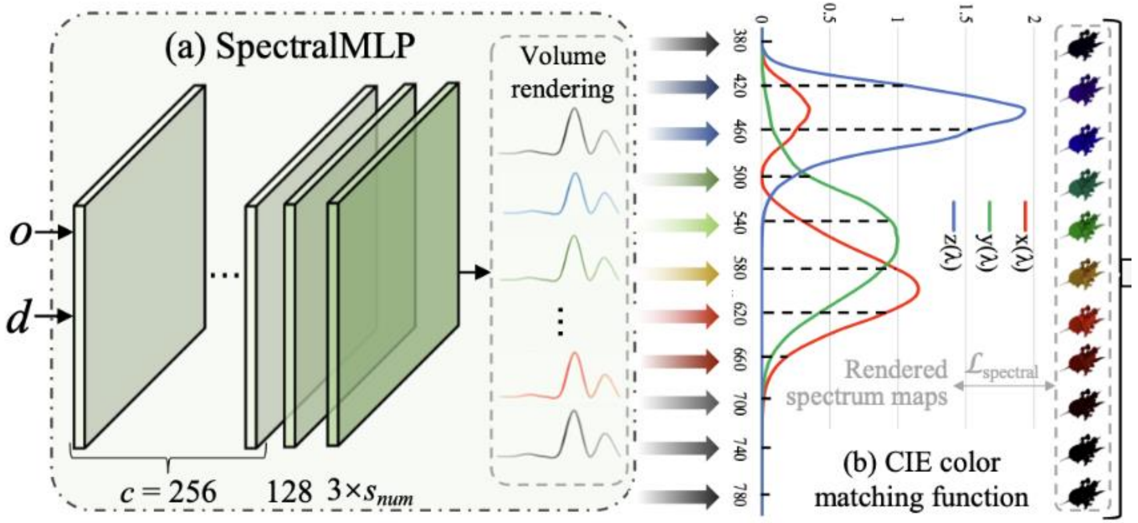

HyperGS [3]

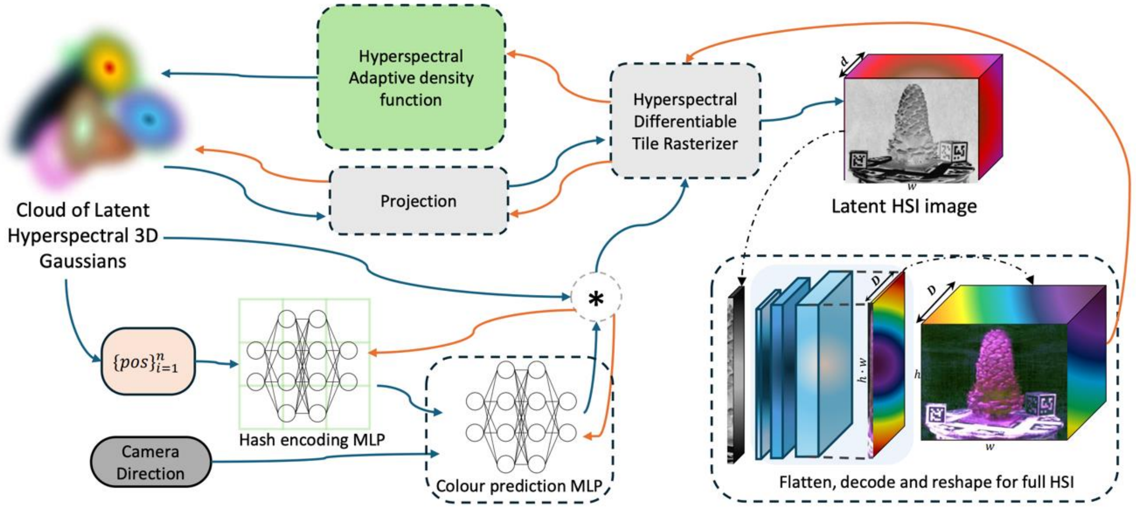

- These approaches do not explicitly leverage the inherent structure of spectral imaging, e.g., **materials have representative spectral signatures**

- None of them allow for material segmentation

[1] Chen, Gerry, et al. "Hyperspectral neural radiance fields." arXiv preprint arXiv:2403.14839 (2024).
[2] Li, Ru, et al. "Spectralnerf: Physically based spectral rendering with neural radiance field." *AAAI 2024.*
[3] Thirgood, Christopher, et al. "Hypergs: Hyperspectral 3d gaussian splatting." Proceedings of the Computer Vision and Pattern Recognition Conference. 2025.

Spectral + Unmixing + 3D = Material-aware NeRF

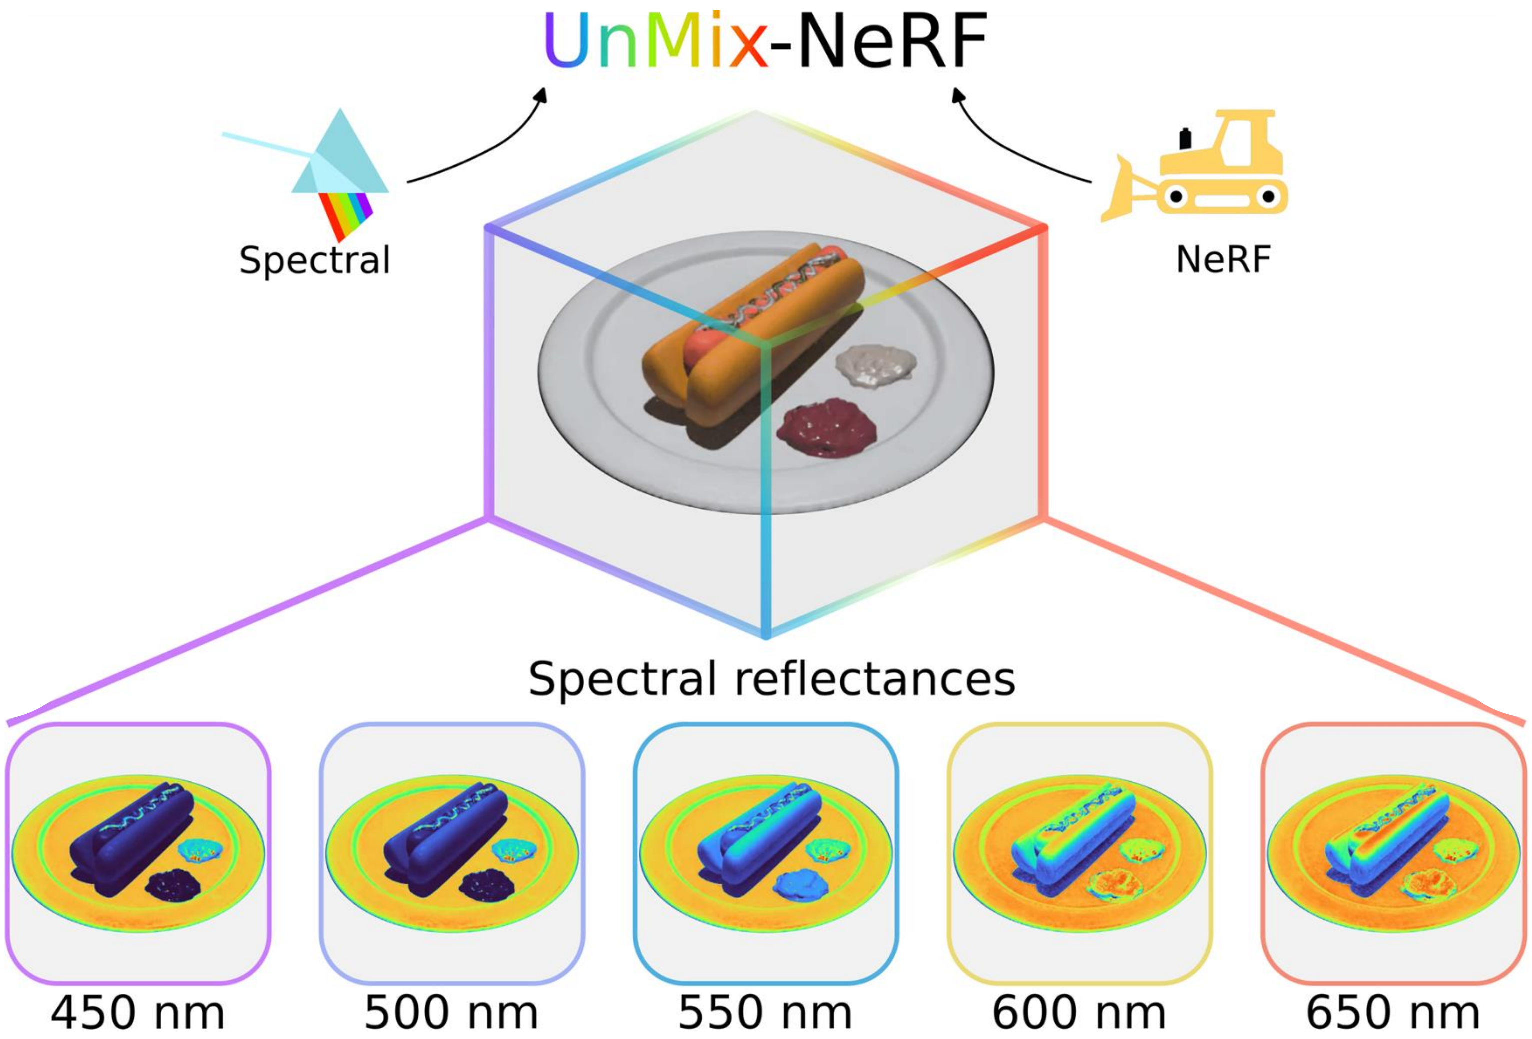

UnMix-NeRF

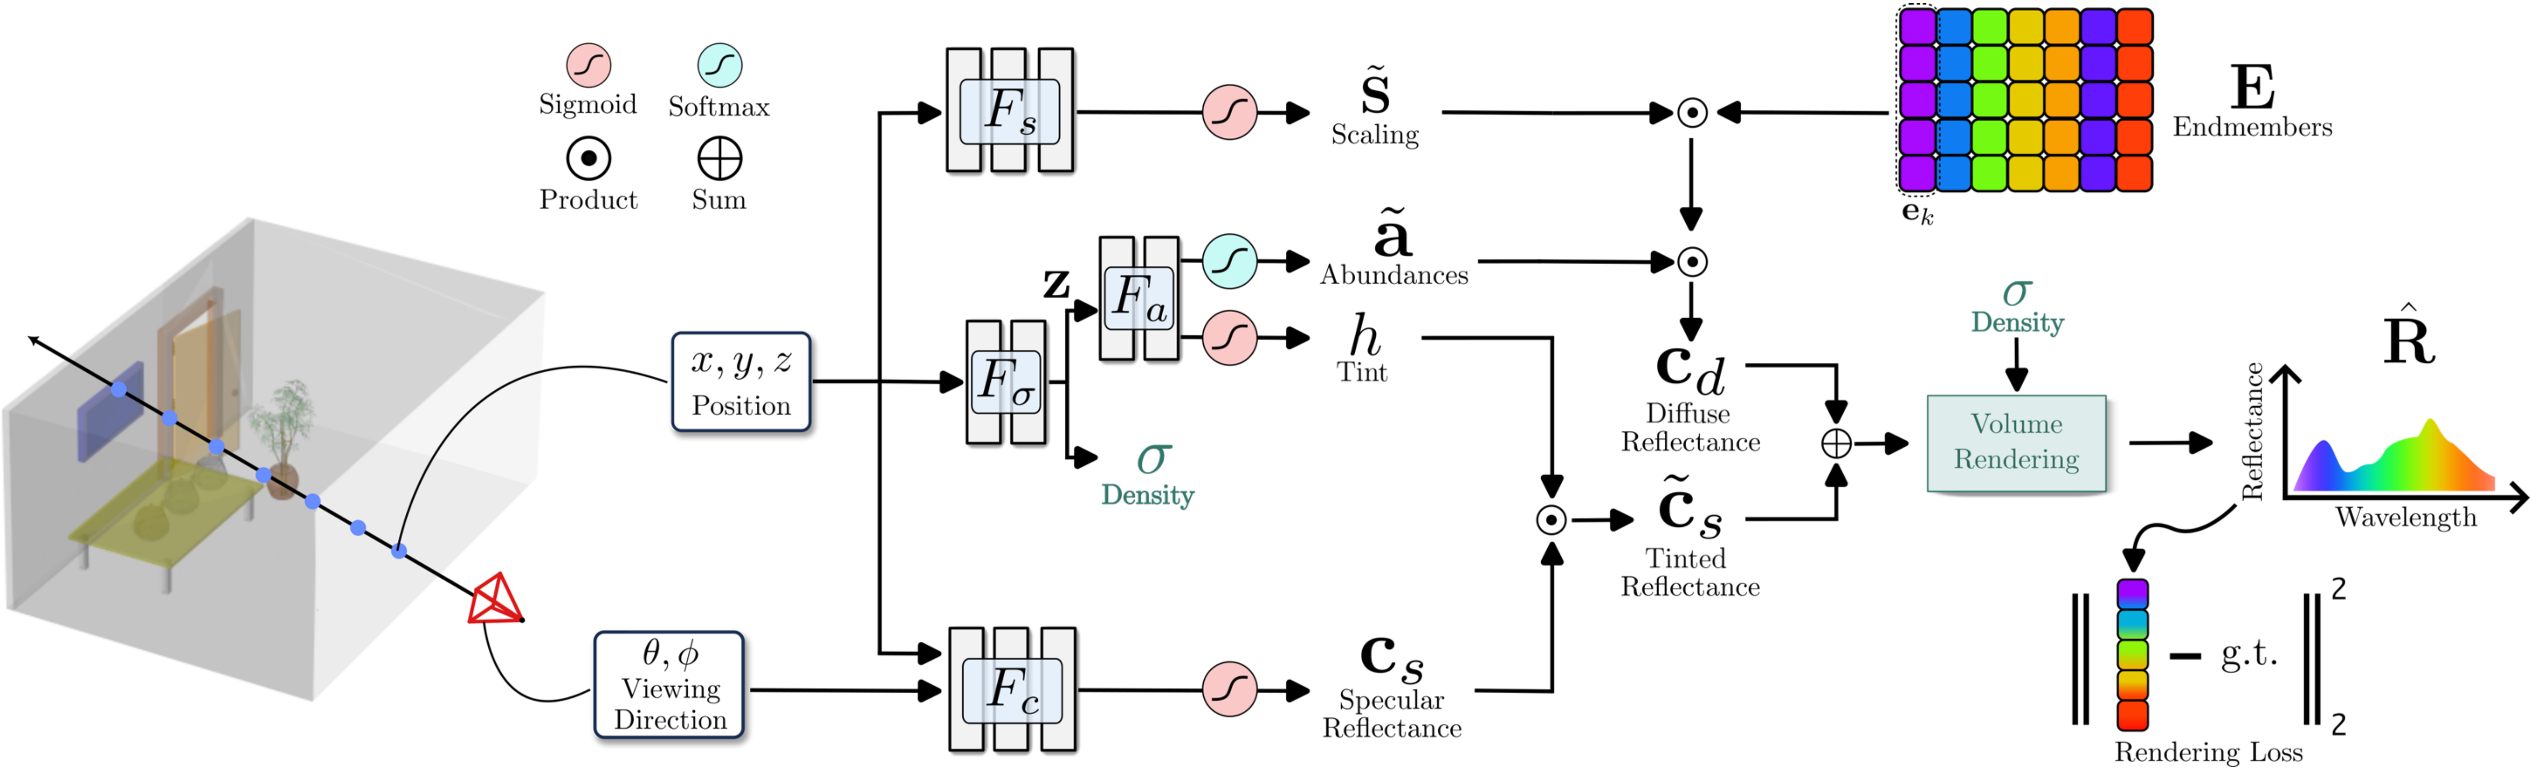

UnMix-NeRF is divided into **diffuse** and **specular** reflectance components.

Diffuse vs Specular reflection

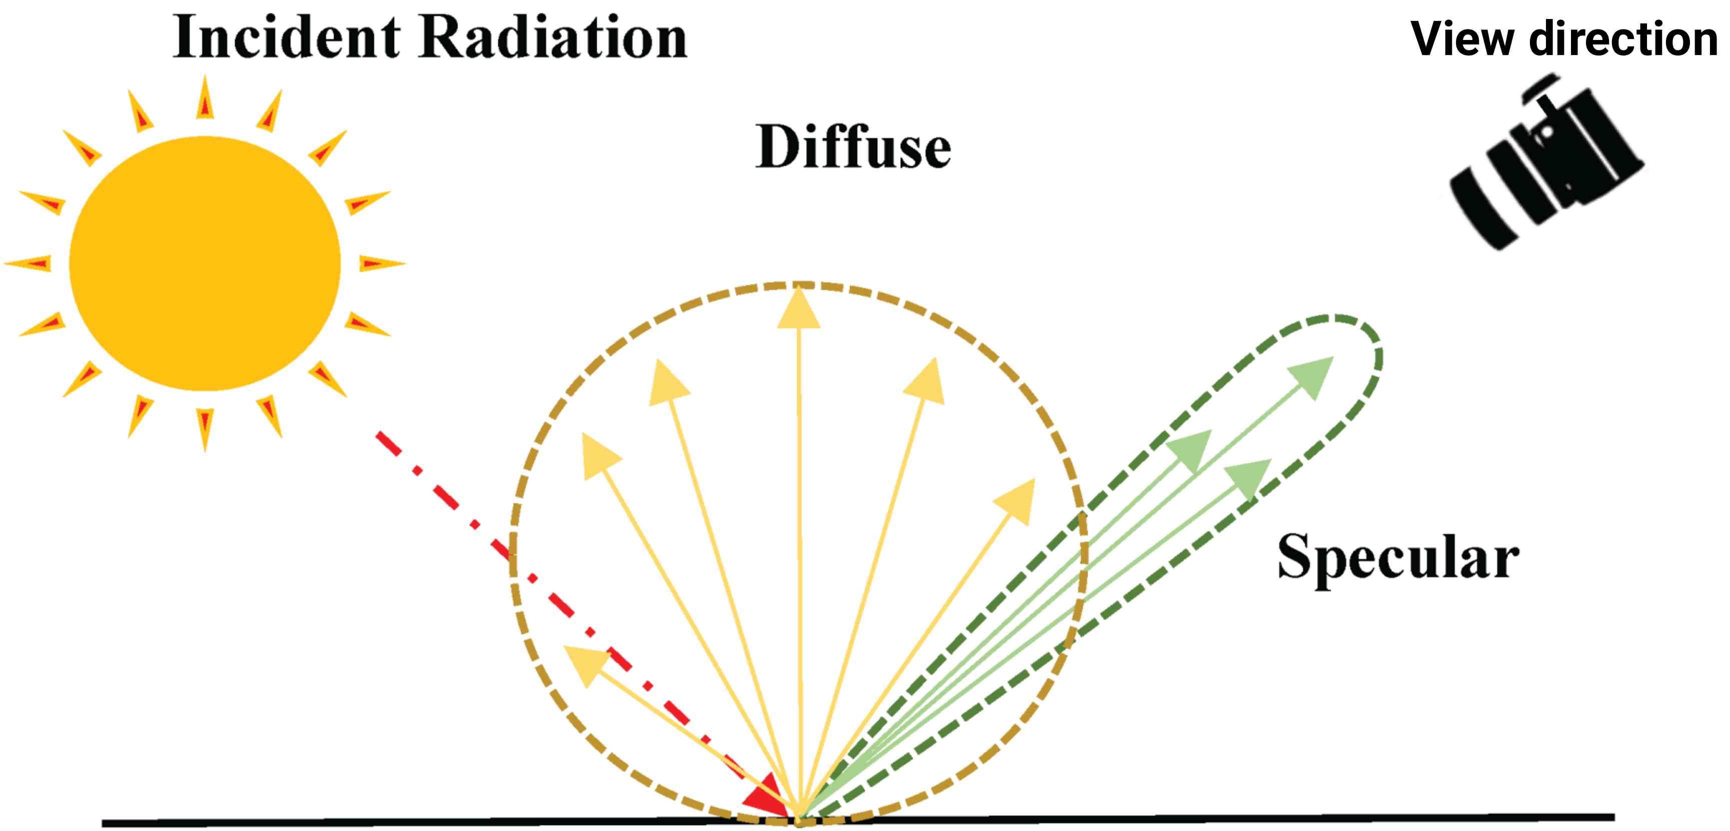

- Diffuse = material-intrinsic, **view-independent**
- Specular = illumination-dependent, **view-dependent**

[1] Jo, Sangho, et al. "S-LIGHT: Synthetic dataset for the separation of diffuse and specular reflection images." Sensors 24.7 (2024): 2286.

UnMix-NeRF

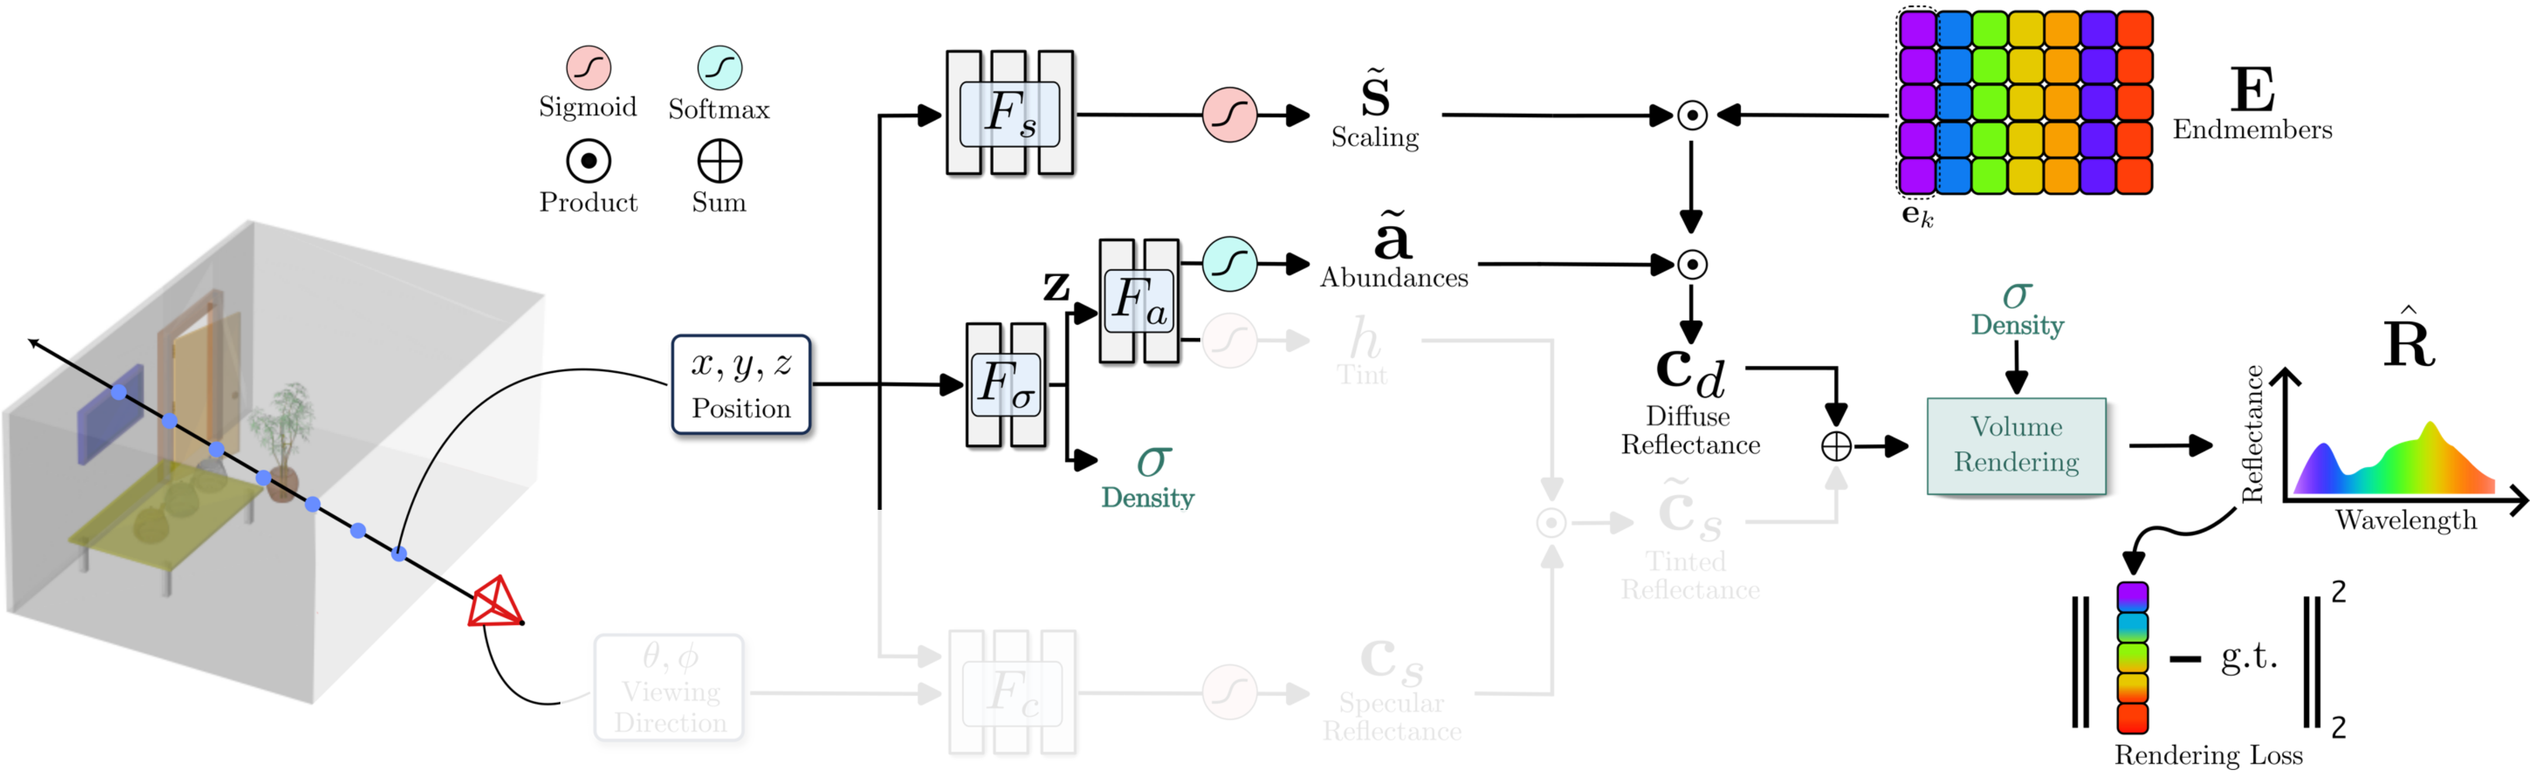

Spectral Unmixing (Diffuse) Field
(view-independent)

Diffuse Reflectance component:

$$\mathbf{c}_d = \mathbf{E}\,\tilde{\mathbf{S}}\tilde{\mathbf{a}}$$

UnMix-NeRF

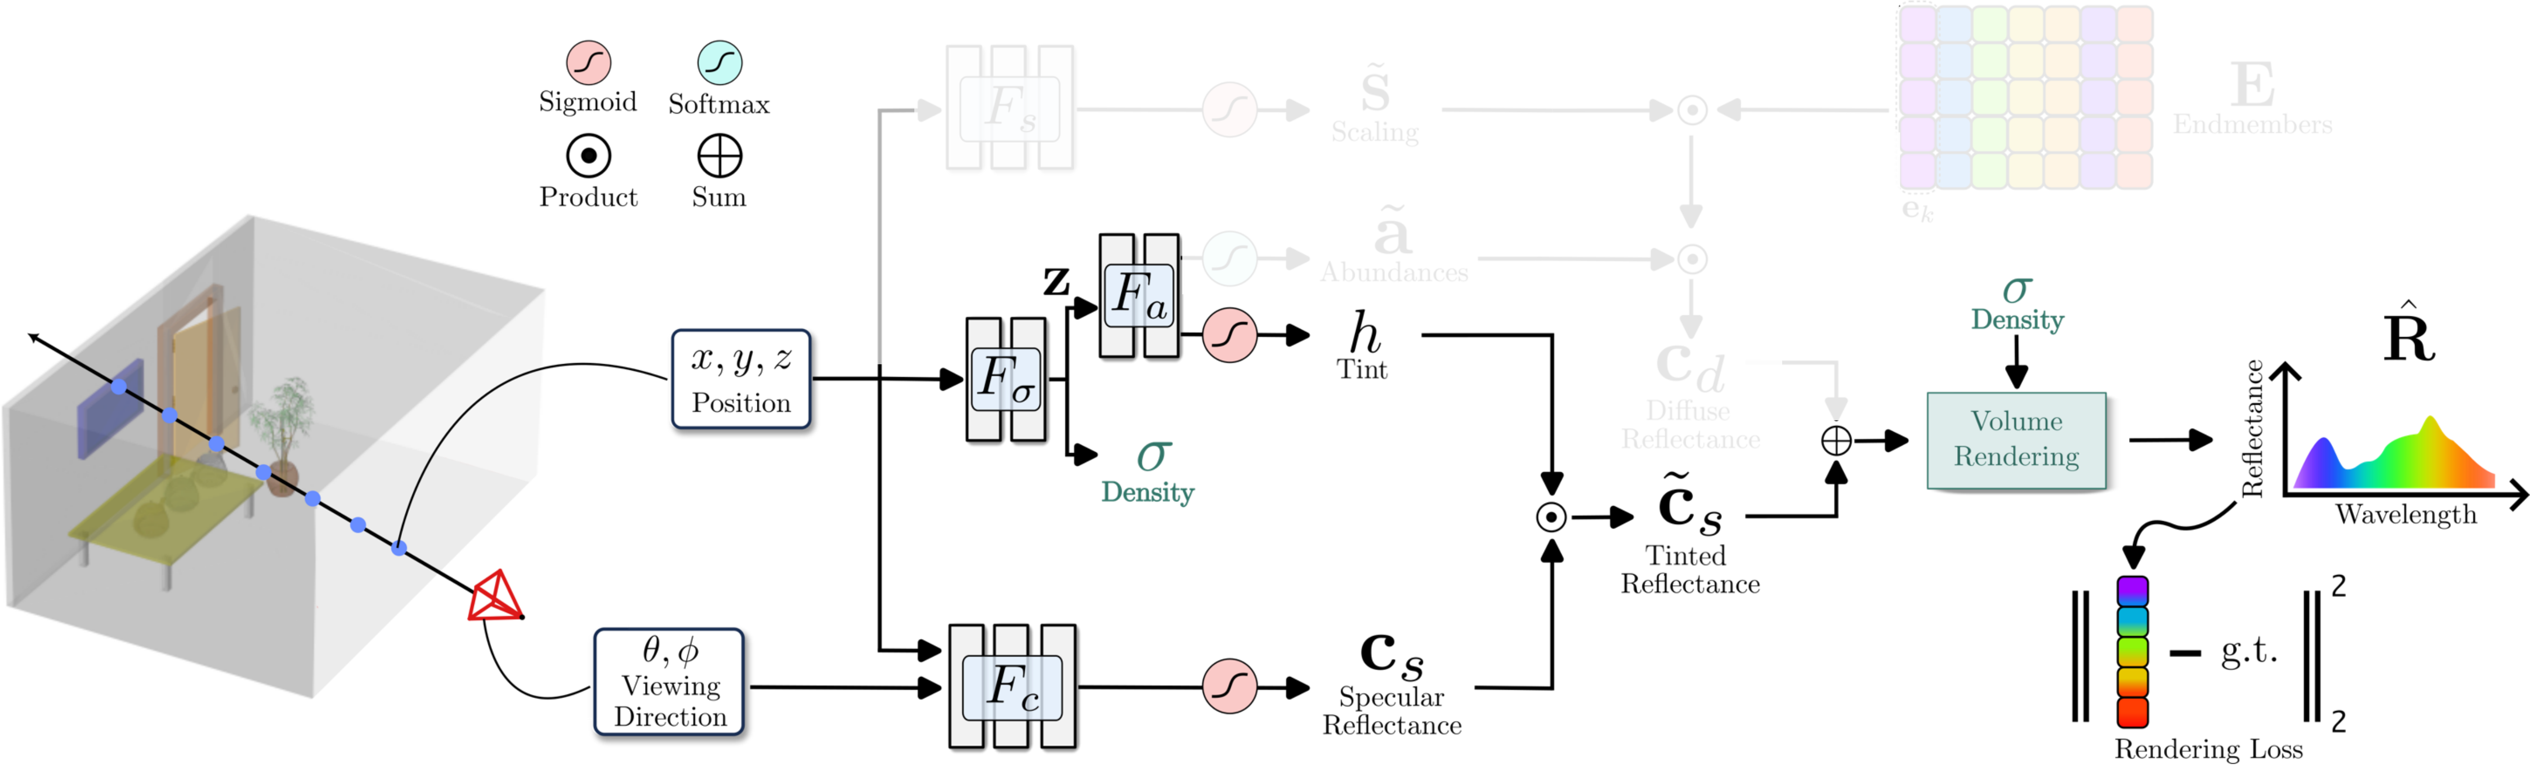

Specular Field
(view-dependent)

Specular Reflectance component:

$$\tilde{\mathbf{c}}_s = h\mathbf{c}_s$$

UnMix-NeRF

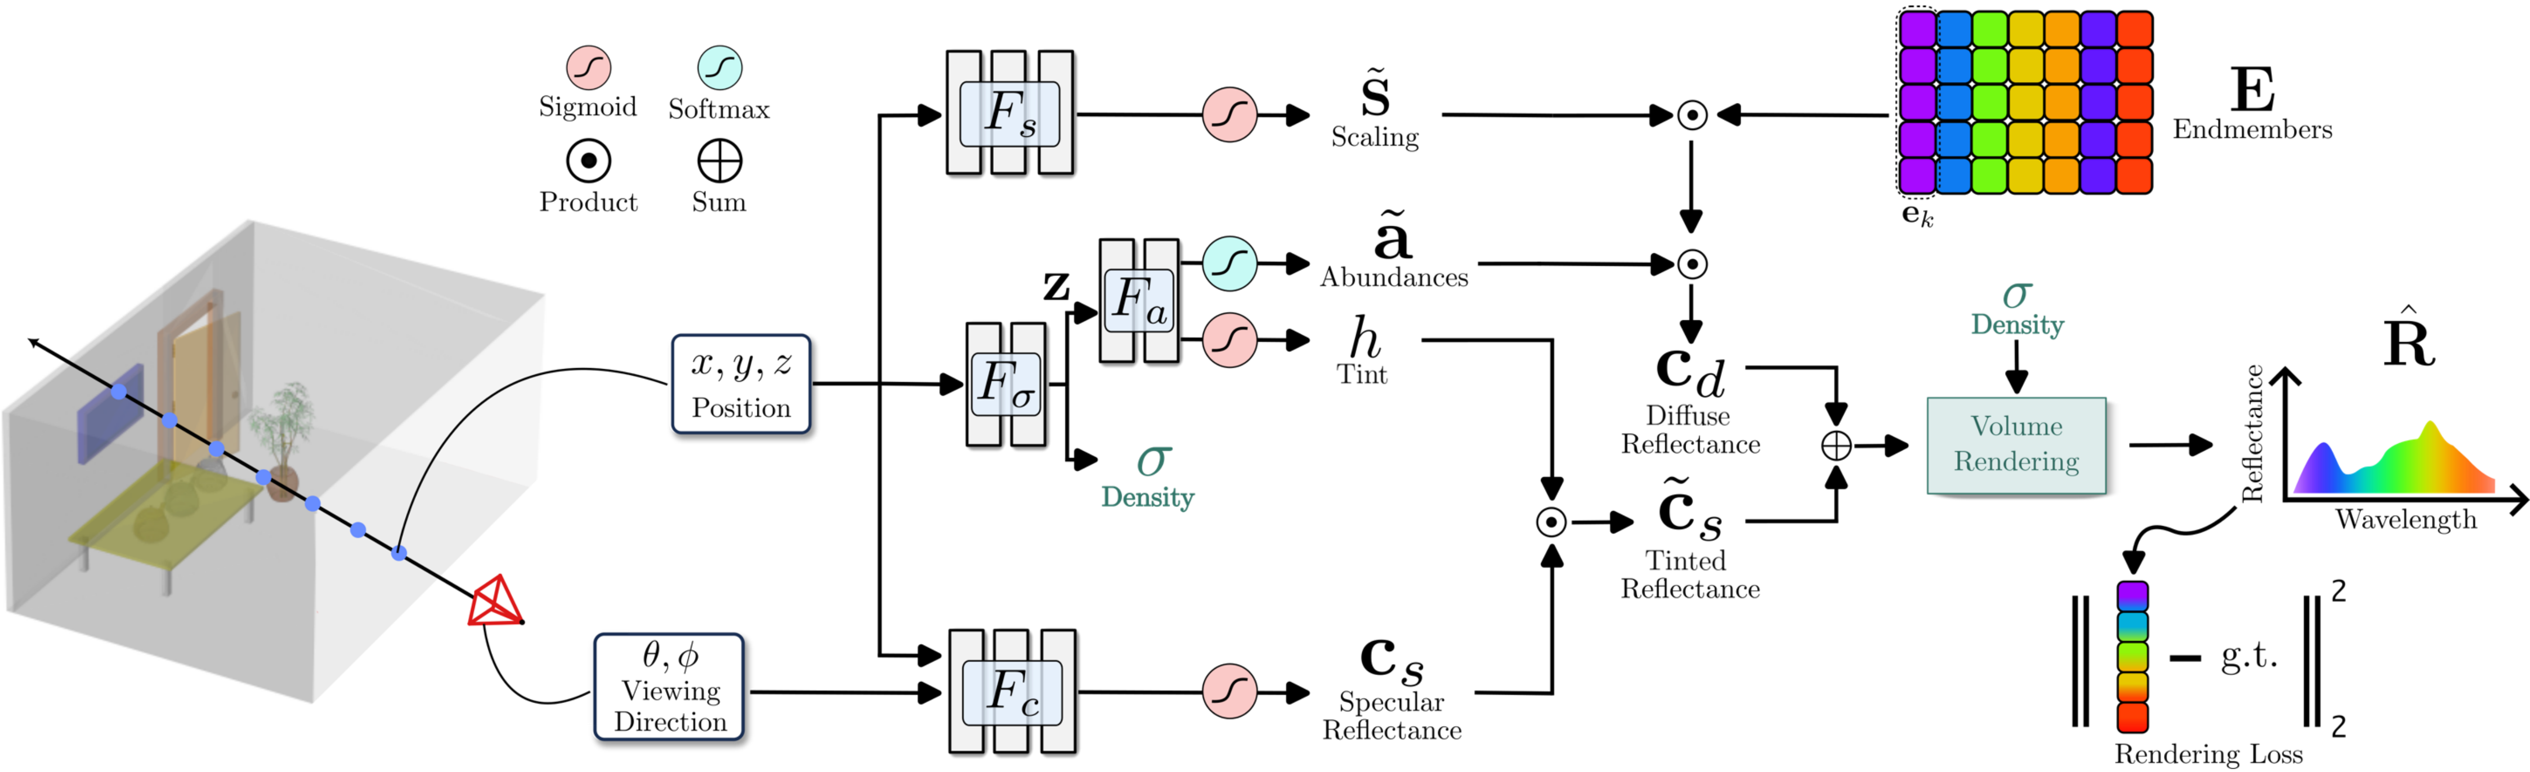

Loss Function

$$L_{\text{spec}} = \sum_{\mathbf{r} \in \mathcal{R}} \|C(\mathbf{r}) - C^*(\mathbf{r})\|_2^2 ,$$

$$L_{\text{rgb}} = \sum_{\mathbf{r} \in \mathcal{R}} \left\|C_{\text{rgb}}(\mathbf{r}) - C^*_{\text{rgb}}(\mathbf{r})\right\|_2^2$$

$$L = \lambda_{\text{spec}} \, L_{\text{spec}} + \lambda_{\text{rgb}} \, L_{\text{rgb}}$$

UnMix-NeRF

Unsupervised Material Segmentation via Cluster Probe

Key idea: Use **spectral unmixing outputs** (abundances and endmembers) to segment materials

Given the learned endmembers:

$$\mathbf{p}(\mathbf{r}) = \frac{\mathbf{E}^\top C(\mathbf{r})}{|\mathbf{E}||C(\mathbf{r})|}$$

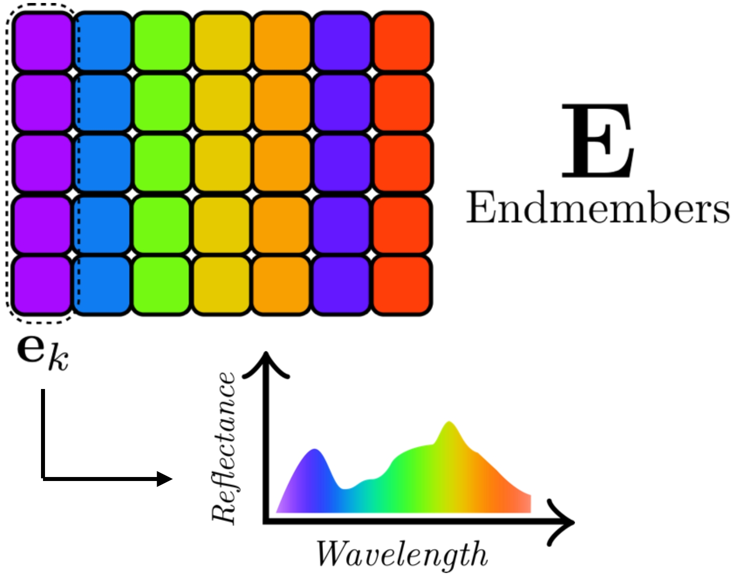

UnMix-NeRF

Unsupervised Material Segmentation via Cluster Probe

Key idea: Use **spectral unmixing outputs** (abundances and endmembers) to segment materials

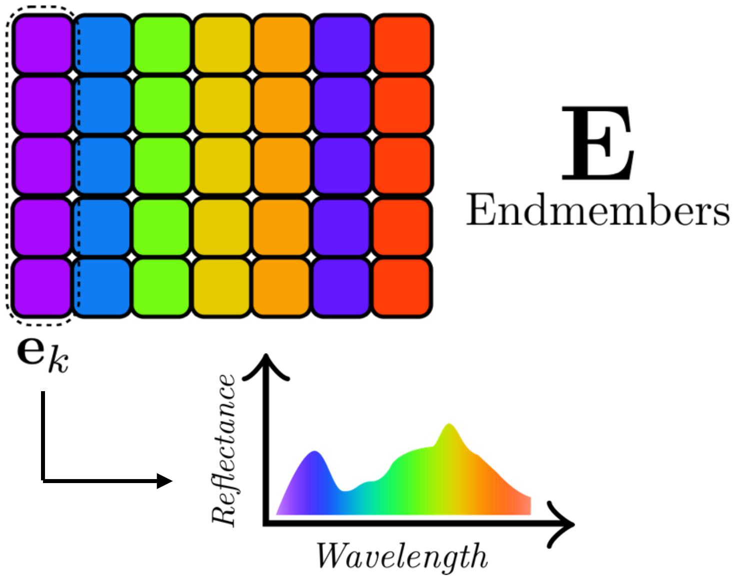

Given the learned endmembers:

$$\mathbf{p}(\mathbf{r}) = \mathrm{softmax}\left(\frac{\mathbf{E}^\top C(\mathbf{r})}{|\mathbf{E}||C(\mathbf{r})|}\right)$$

Each ray is assigned to the material with the highest probability:

$$m(\mathbf{r}) = \arg\max_k p_k(\mathbf{r})$$

No manual labels needed, **segmentation emerges from spectral structure**

Experimental Setup: Datasets

NeSpoF

- 21 spectral bands
- 450nm to 560 nm
- 4 scenes
- Synthetic dataset

Surface Optics

- 128 spectral bands
- 370nm to 1100 nm
- 4 scenes
- Real dataset

BaySpec

- 140 spectral bands
- 400nm to 1110 nm
- 3 scenes
- Real dataset

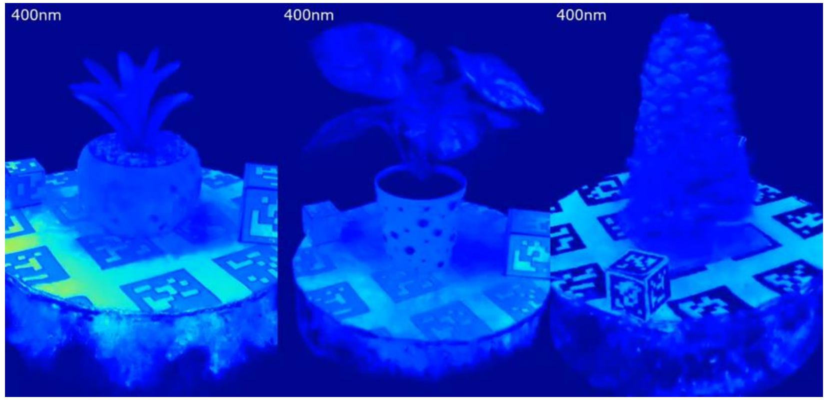

Experimental Setup

NeSpoF

Surface Optics

BaySpec

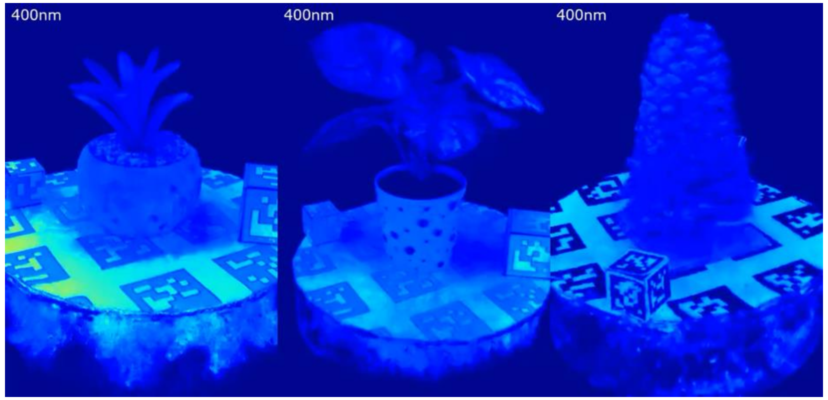

We extended the NeSpoF dataset by providing ground-truth material labels for all the synthetic scenes

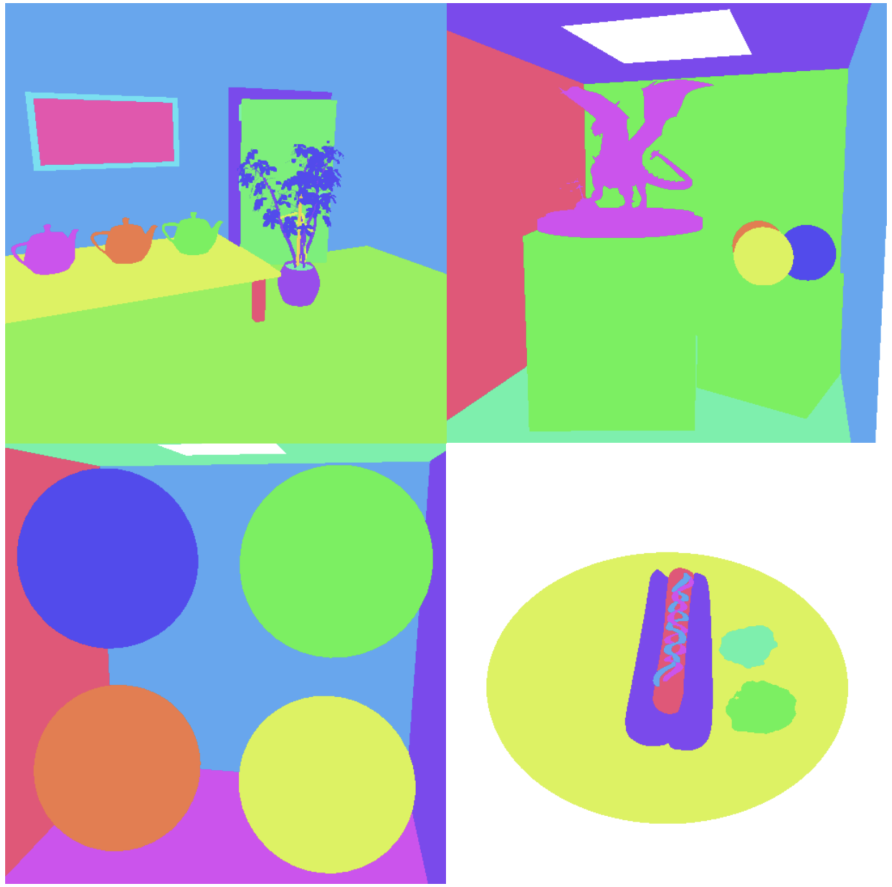

Quantitative Results: NeSpoF

Hotdog

Cbox-sphere

Cbox-dragon

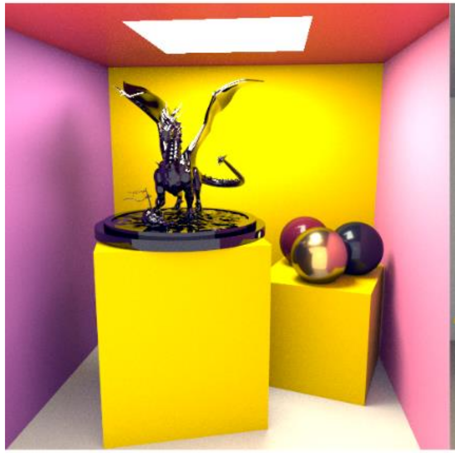

Ajar

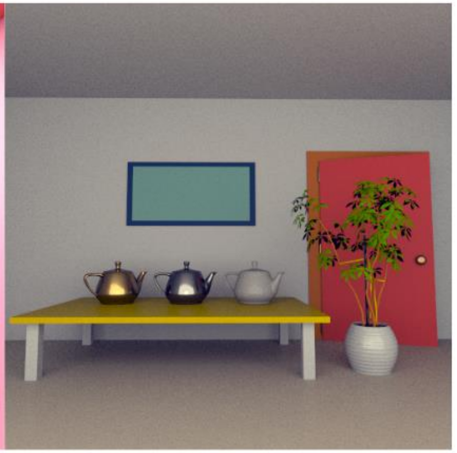

Method	Scene	PSNR ↑	RMSE ↓	Time
HS-NeRF*	avg.	26.0	0.04	5 hours
NeSpoF	avg.	33.0	0.02	11.9 hours
Ours	ajar	38.09	0.01	43 min
	hotdog	34.47	0.01	45 min
	cbox-dragon	32.21	0.02	45 min
	cbox-sphere	27.96	0.03	44 min
	avg.	**33.2**	**0.02**	**44 min**

Table 2. Comparison on the NeSpoF dataset. HS-NeRF* and NeSpoF report averaged results (avg.) over 4 scenes; for Ours, per-scene results and the overall averaged metrics are provided.

Quantitative Results: Surface Optics

Method	Rosemary				Basil			
	PSNR ↑	SSIM↑	SAM↓	RMSE↓	PSNR ↑	SSIM↑	SAM↓	RMSE↓
NeRF	8.42	0.7461	0.0284	0.3560	9.91	0.5534	0.0769	0.5256
MipNeRF	13.64*	0.5684*	1000*	0.2083*	10.11	0.5878	0.0728	0.5334
TensoRF	12.1	0.73351	0.0212	0.2662	15.23	0.5811	0.0435	0.3628
Nerfacto	18.66	0.8836	0.0078	0.1205	16.54	0.7915	0.0176	0.1655
MipNerf360	8.47	0.7518	0.0876	0.3825	13.92	0.8584	0.0497	0.2035
HS-NeRF	*18.60	*0.887	*0.0077	*0.1187	*16.81	*0.771	*0.0172	*0.1587
3DGS	25.56	*0.9695*	0.0028	0.0534	21.19	0.9385	0.0101	0.0897
HyperGS	*26.77*	**0.9845**	*0.0021*	*0.0445*	*25.30*	*0.9503*	*0.00514*	*0.0569*
Ours	**28.91**	0.9355	**0.0019**	**0.0332**	**29.21**	**0.9584**	**0.0043**	**0.0364**

Best on Rosemary and Basil (+2.14/+3.91 PSNR vs. HyperGS) with lowest SAM and RMSE, confirming accurate spectral reconstruction across diverse materials

Rosemary Scene

Basil Scene

Quantitative Results: BaySpec

Method	Pinecone				Caladium				Anacampseros			
	PSNR ↑	SSIM↑	SAM↓	RMSE↓	PSNR ↑	SSIM↑	SAM↓	RMSE↓	PSNR ↑	SSIM↑	SAM↓	RMSE↓
NeRF	22.82	0.6113	0.0446	0.0728	23.12	0.58348	0.0491	0.0709	24.12	0.6220	0.0384	0.0623
MipNeRF	21.45	0.5738	0.0410	0.0856	23.36	0.5935	0.0487	0.0685	23.43	0.6160	0.0408	0.0786
TensoRF	24.12	0.6454	0.0593	0.0625	24.79	0.6424	0.0516	0.0577	25.07	0.6569	0.0394	0.0558
Nerfacto	15.36	0.4935	0.0707	0.1709	20.67	0.6208	0.0529	0.0945	21.32	0.6423	0.0417	0.0867
MipNeRF360	20.93	0.7355	0.0279	0.0507	26.93	0.7371	0.0332	0.0461	26.73	0.7601	0.0230	0.0461
HS-NeRF	20.07	0.581	0.0725	0.1521	19.084	0.705	0.0533	0.0902	20.32	0.7260	0.0345	0.0789
3DGS	22.65	0.6039	0.0668	0.0819	23.50	0.7131	0.2889	0.0758	22.59	0.5786	0.0447	0.0853
HyperGS	*27.0*	*0.7509*	*0.0309*	*0.0447*	*27.70*	*0.8354*	*0.0271*	*0.0414*	*26.62*	*0.7545*	*0.0183*	*0.0460*
Ours	**27.13**	**0.8174**	**0.0287**	**0.0429**	**30.08**	**0.8541**	**0.0237**	**0.0312**	**28.20**	**0.7612**	**0.0154**	**0.0392**

Consistently best across scenes, highest PSNR/SSIM and lowest SAM/RMSE. Explicitly modeling spectral unmixing in NeRF yields more accurate spectra than per-pixel regression

Anacampseros Caladium Pinecone

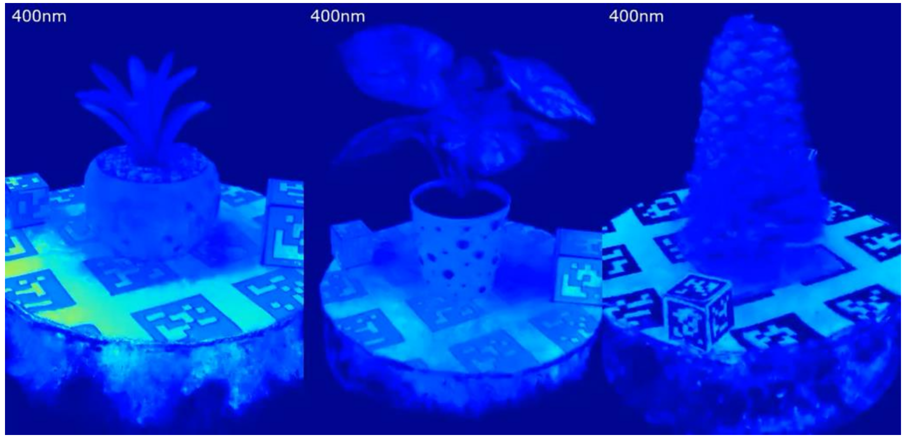

Qualitative Results: BaySpec

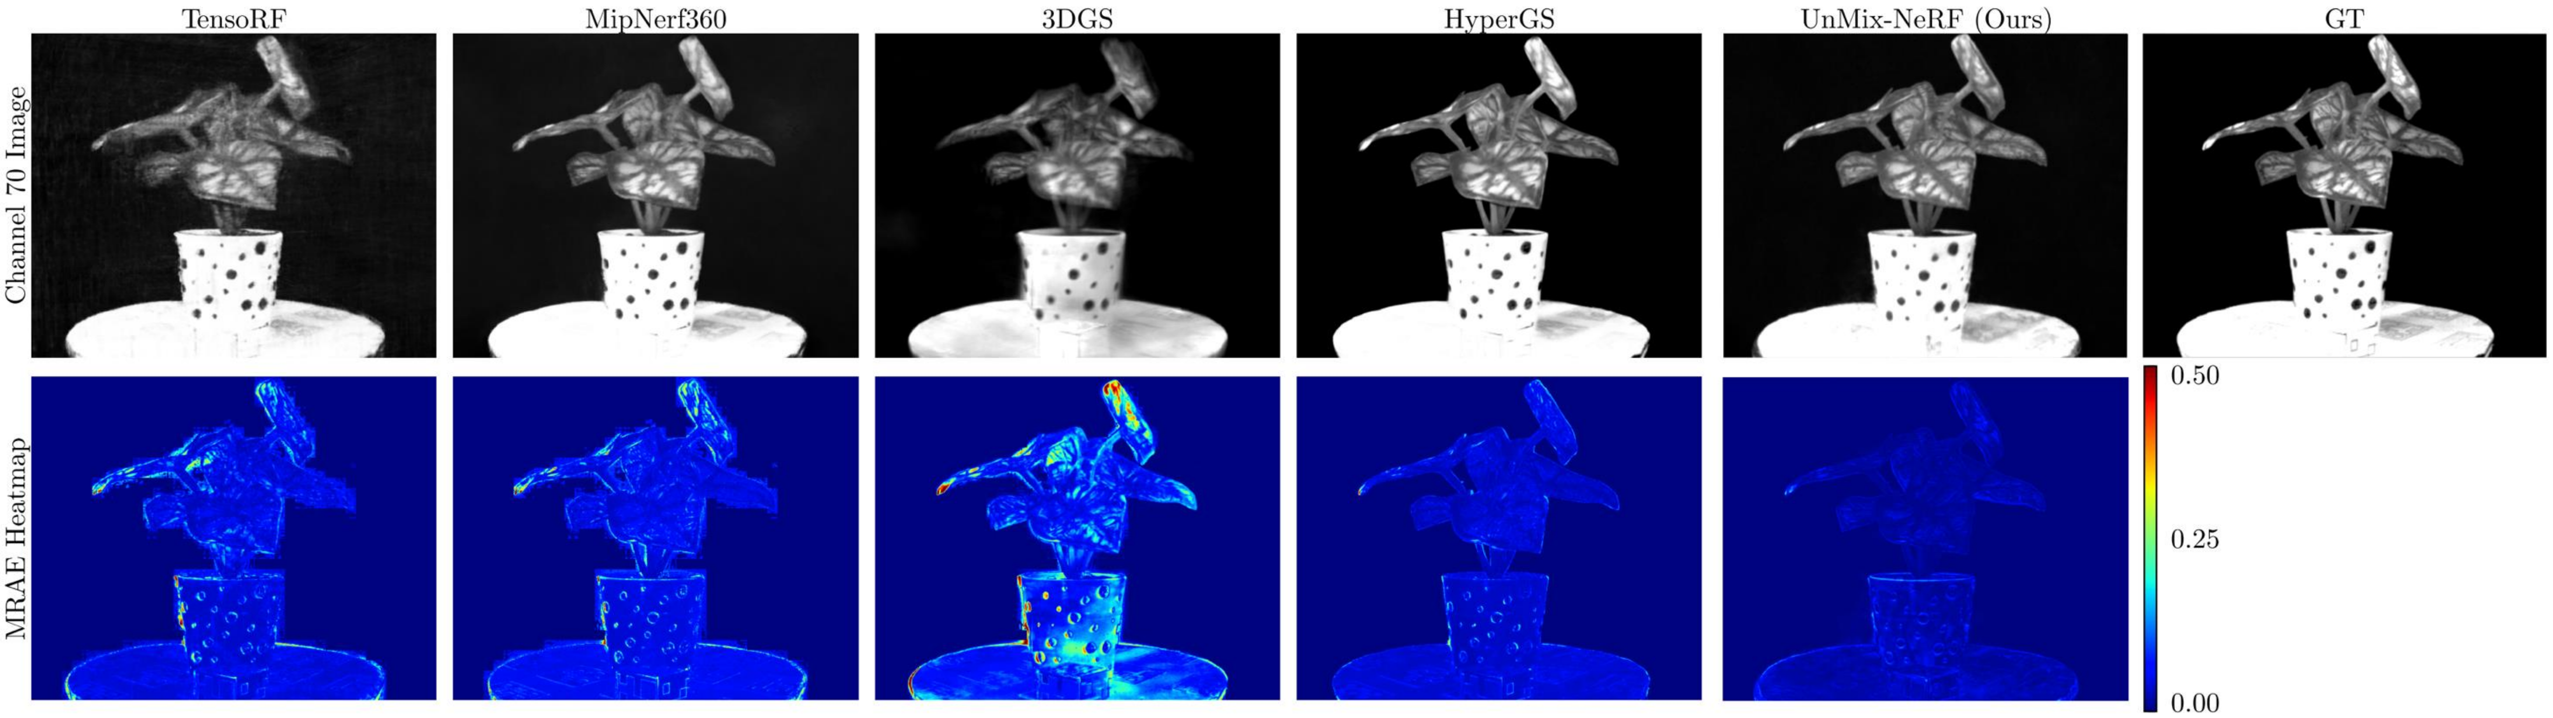

Frame number 51 in the Caladium scene

Our method achieves the most accurate spectral predictions, significantly reducing reconstruction artifacts and preserving fine-grained spectral details

Qualitative Results: Hotdog

Spectral Reflectances

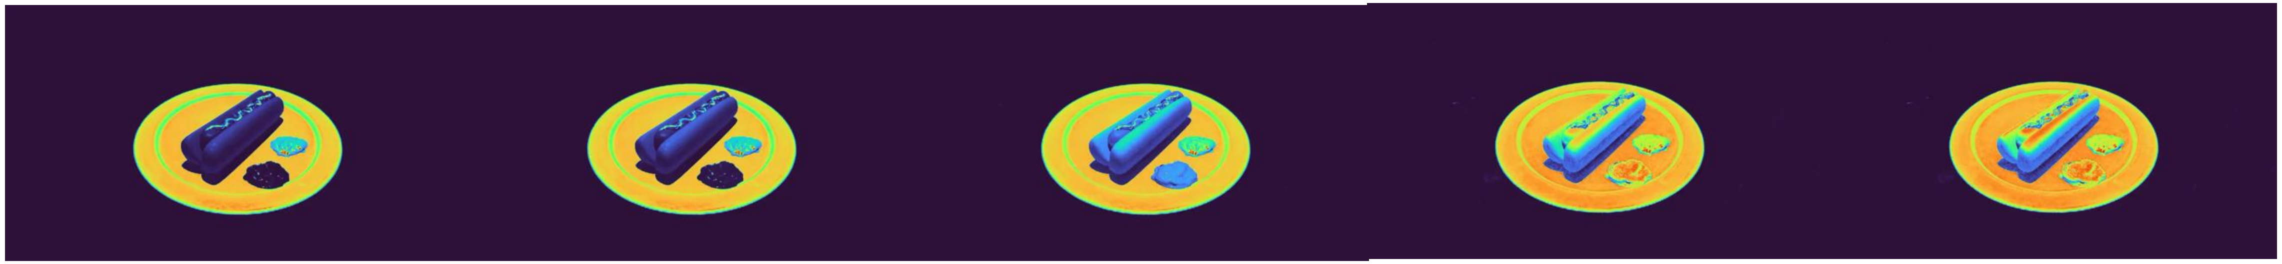

450 nm | 500 nm | 550 nm | 600 nm | 650 nm

Specular Reflectance

Learned Material Abundances

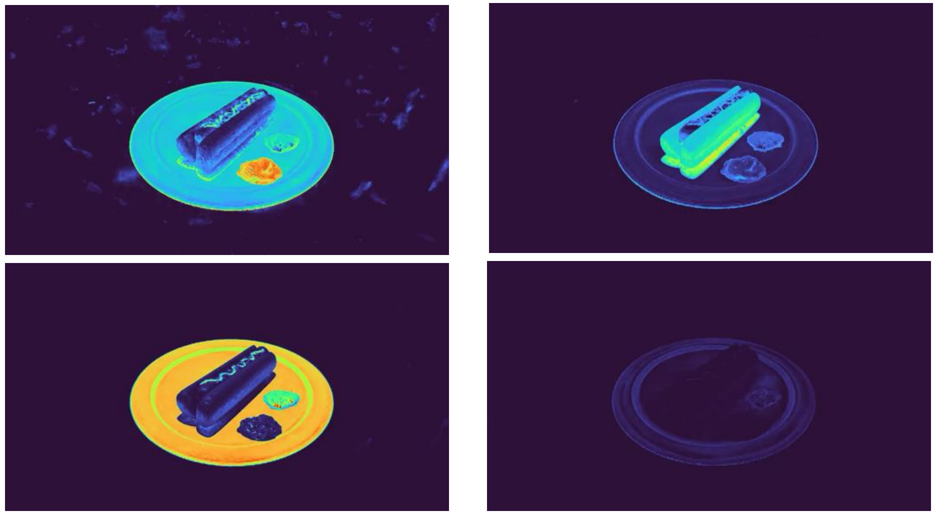

Qualitative Results: Hotdog

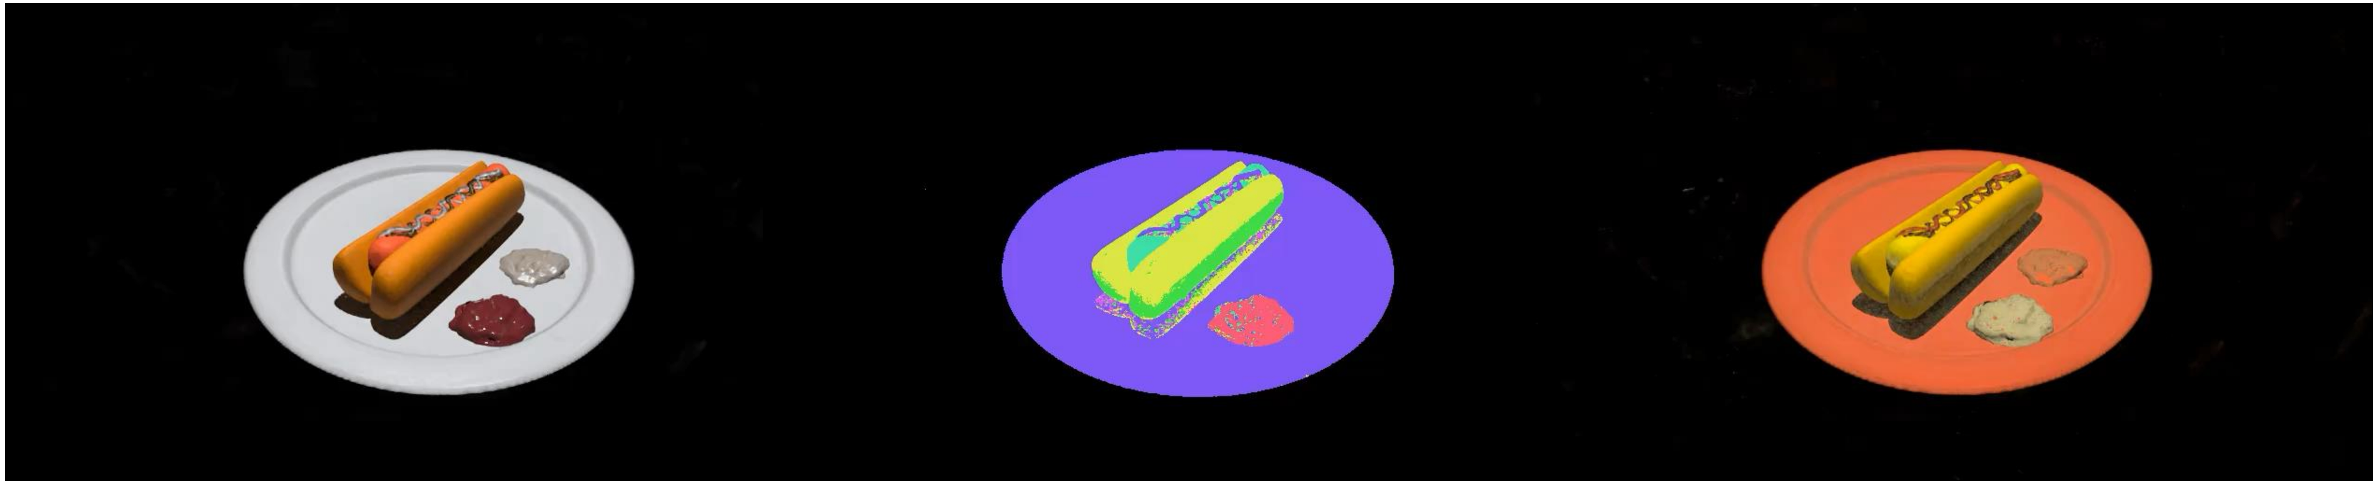

RGB | Unsupervised Material Segmentation | Scene Editing

Qualitative Results: Ajar

Visualization of Learned Material Abundances

RGB

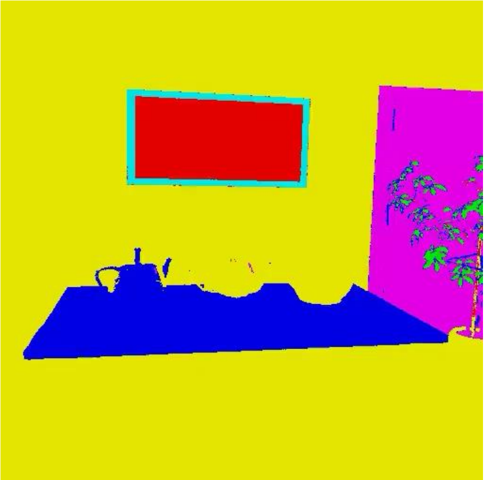

Unsupervised Material Segmentation

PCA Visualization

Take Aways

UnMix-NeRF is the first 3D NeRF to integrate spectral unmixing, jointly learning endmembers and abundances for material-aware novel view synthesis

Spectral segmentation emerges from endmember estimation, materials separate via an unsupervised cluster probe, with no labels

Endmember-level editing enables intuitive scene control, letting you swap/retint materials consistently across all novel views

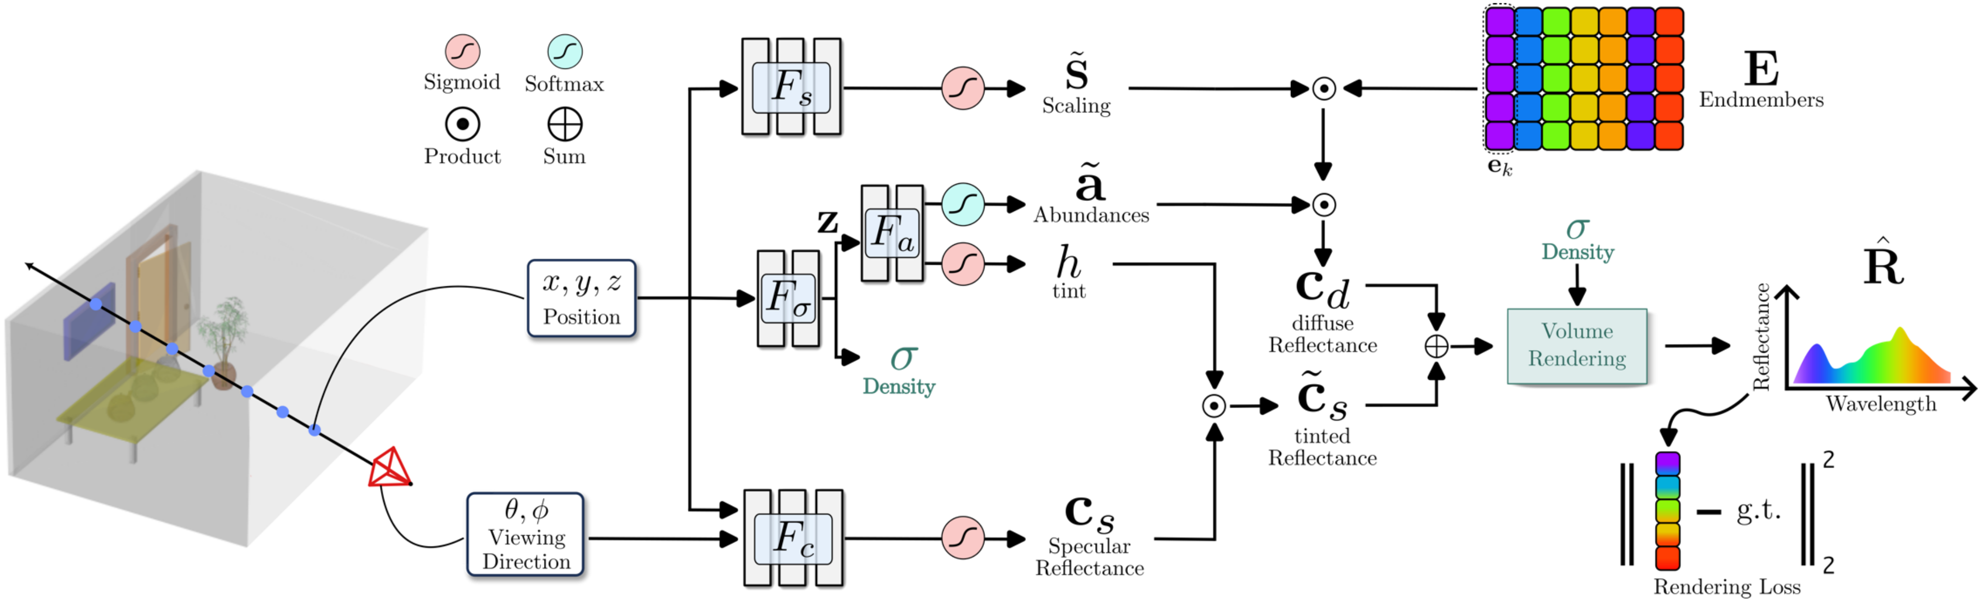

Thank You!

Project Page

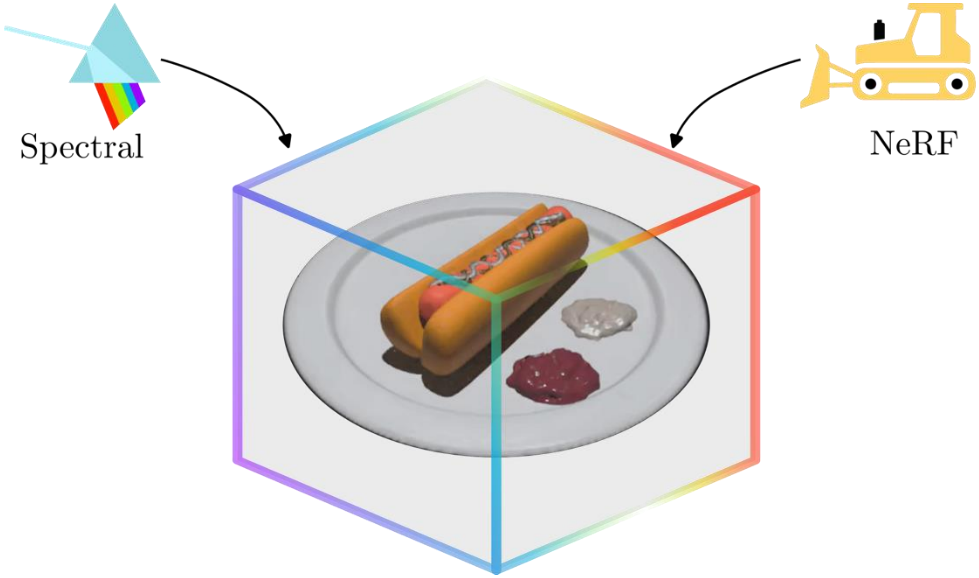

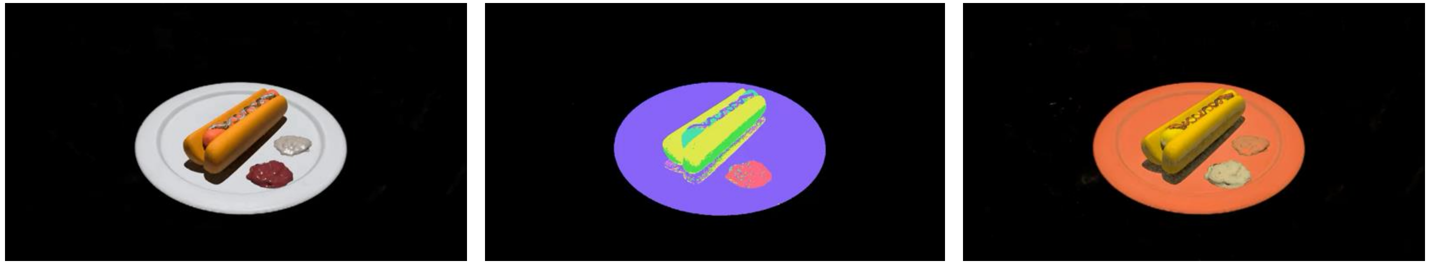

Visit us at **poster #1682**,
Thu, Oct 23rd