

UnMix-NeRF: Spectral Unmixing Meets Neural Radiance Fields

Supplementary Material

7. Network

Our proposed method, UnMix-NeRF, builds upon Nerfacto by incorporating spectral unmixing into the radiance field representation. The full architecture is detailed in Table 5, following the structure of Figure 2 in the main paper. The *Base Geometry MLP* processes the 12D NeRF positional encoding and outputs density σ along with a 15D geometry feature vector. This feature vector, along with the positional encoding, is then passed to the *Feature MLP*, which predicts per-point material abundances and the specular tint.

To allow for per-point spectral scaling, the *Scaling MLP* separately processes the position encoding and predicts K per-material scaling factors, ensuring a flexible spectral reconstruction. To handle view-dependent effects, the *Directional MLP* receives the direction encoding (16D SH encoding) concatenated with the positional encoding (12D), predicting a per-wavelength specular reflectance using a final sigmoid activation to ensure physically plausible outputs.

This modular design ensures that spectral unmixing is seamlessly integrated into the NeRF volumetric rendering process, allowing for accurate hyperspectral novel view synthesis while also enabling material segmentation through abundance-based clustering.

Module	Input Features	Output Features
Base Geometry MLP		
Layer 1: Linear + ReLU	12 (pos. encoding)	64
Layer 2: Linear + ReLU	64	64
Layer 3: Linear + ReLU	64	64
Layer 4: Linear + ReLU (skip)	64 + 12	64
Layer 5: Linear	64	16 (density and features)
Feature MLP (Abundances + Specular Tint)		
Layer 1: Linear + ReLU	27 (pos. enc. 12 + geo. feat 15)	64
Layer 2: Linear + ReLU	64	64
Layer 3: Linear	64	$K + 1$ (abundances and tint)
Scaling MLP (Endmember Scalars)		
Layer 1: Linear + ReLU	12 (pos. enc.)	64
Layer 2: Linear + ReLU	64	64
Layer 3: Linear	64	K
Directional MLP (Specular Reflectance)		
Layer 1: Linear + ReLU	28 (dir. enc. 16 + pos. enc. 12)	16
Layer 2: Linear + Sigmoid	16	Wavelengths (e.g. 128)

Table 5. **UnMix-NeRF architecture.** Our method extends Nerfacto by integrating spectral unmixing into the NeRF formulation.

8. Additionally HS-NeRF results

We present the quantitative results for the Tools and Origami scenes from the Surface Optics dataset. Unlike the Rosemary and Basil scenes, these cases exhibit significant convergence issues due to corrupted original pose files. Consequently, the optimization process fails to achieve

high-quality reconstructions, leading to suboptimal qualitative outcomes. Moreover, since 3DGS and HyperGS were evaluated using different, uncorrupted pose files, any direct comparison under these conditions would be inherently unfair. Therefore, we exclude the results for 3DGS and HyperGS on these scenes to maintain a consistent and fair evaluation. We reached out to the HyperGS authors, but they did not provide their pose files.

Method	Tools				Origami			
	PSNR \uparrow	SSIM \uparrow	SAM \downarrow	RMSE \downarrow	PSNR \uparrow	SSIM \uparrow	SAM \downarrow	RMSE \downarrow
NeRF	11.61	0.4962	0.0610	0.3018	13.64	0.5684	0.0835	0.2083
MipNeRF	12.78	0.5213	0.0598	0.2781	11.697	0.5149	0.0956	0.2595
TensoRF	11.697*	0.5149*	0.0956*	0.2595*	12.98	0.4488	0.0776	0.2314
Nerfacto	16.254	0.6135	0.0198	0.1549	14.02	0.5028	0.0953	0.1993
MipNerf360	16.80	0.7241	0.0832	0.1482	9.93	0.3951	0.3271	0.3288
HS-NeRF	*12.001	*0.355	*0.470	*0.185	10.359	0.4530	0.3197	0.3188
Ours	17.347	0.4729	0.0174	0.1357	15.973	0.3251	0.086	0.1403

Table 6. Quantitative results on the Surface Optics dataset for the *Tools* and *Origami* scenes.

9. Extended NeSpoF synthetic dataset

To advance material segmentation in multi-view settings, we extend the NeSpoF dataset by providing ground-truth material labels for all scenes in the synthetic scenes. These labels are available for every viewpoint, enabling precise evaluation of material segmentation methods. We generate the material annotations by rendering the corresponding material index for each object in the scene, ensuring consistency across views. This dataset extension will be publicly released as a benchmark for evaluating material segmentation. Figure 6 showcases the rendered material annotations for various synthetic scenes, where distinct colors correspond to different materials.




Figure 6. **Extended NeSpoF Synthetic Dataset.** ground-truth material segmentation for different synthetic scenes.

10. Limitations

The number of endmembers is a key hyperparameter in spectral unmixing, determined by scene complexity and material diversity. Outdoor scenes typically require 3–5 endmembers (e.g., vegetation, soil, concrete) while more diverse environments up to 10 or more. This number can be estimated from known materials or through methods such as VCA. The theoretical upper bound is the number of spectral channels plus one. A potential inherited limitation of our work is its sensitivity to the incorrect estimation of the number of learnable endmembers, which can affect both reconstruction quality and interpretability as shown in the following image. Also, if predicted endmembers are not distinct enough, it may lead to metamerism.



Figure 7. **Effect of k (number of endmembers) on reconstruction.** With $k=1$, the model collapses to a near-uniform image (underfitting). As k increases ($3 \rightarrow 7$), edges, textures, and color fidelity progressively improve, approaching the ground truth (GT). Small k fails to separate materials; larger k captures spectral variability better and yields higher-fidelity reconstructions.