

UnMix-NeRF: Spectral Unmixing Meets Neural Radiance Fields

Fabian Perez^{1,2,†}, Sara Rojas², Carlos Hinojosa^{2,*}, Hoover Rueda-Chacón¹, Bernard Ghanem²
¹Universidad Industrial de Santander ²KAUST




Figure 1. UnMix-NeRF: A hyperspectral novel view synthesis framework that leverages spectral unmixing for scene editing and unsupervised material segmentation. By exploiting endmembers and abundances, our approach enables material separation and scene manipulation.

Abstract

Neural Radiance Field (NeRF)-based segmentation methods focus on object semantics and rely solely on RGB data, lacking intrinsic material properties. This limitation restricts accurate material perception, which is crucial for robotics, augmented reality, simulation, and other applications. We introduce UnMix-NeRF, a framework that integrates spectral unmixing into NeRF, enabling joint hyperspectral novel view synthesis and unsupervised material segmentation. Our method models spectral reflectance via diffuse and specular components, where a learned dictionary of global endmembers represents pure material signatures, and per-point abundances capture their distribution. For material segmentation, we use spectral signature predictions along learned endmembers, allowing unsupervised material clustering. Additionally, UnMix-NeRF enables scene editing by modifying learned endmember dictionaries for flexible material-based appearance manipulation. Extensive experiments validate our approach, demonstrating superior spectral reconstruction and material segmentation to existing methods. Project page: <https://www.factral.co/UnMix-NeRF>.

1. Introduction

Neural Radiance Fields (NeRF) [31] have revolutionized 3D scene representation, enabling photorealistic novel view synthesis through implicit volumetric rendering. Recent advances have extended NeRF to other tasks such as segmentation [7], object detection [14], and spatial reasoning [16], leveraging pre-trained vision models such as SAM [19] and CLIP [36] for higher-level scene understanding. Among these, semantic segmentation has emerged as a powerful tool, enabling the grouping of regions into object-level categories in multi-view settings [30]. A more challenging task lies in understanding and segmenting materials, moving beyond recognizing what an object is to understanding what it is composed of. Materials are fundamental to how we interpret and interact with the world, as they define surface properties, functionality, and usability, enabling richer scene understanding and making them essential for applications such as robot autonomy [1], augmented reality [48], and simulation [41], where accurate material perception is crucial for decision-making and interaction.

Most existing 3D segmentation approaches focus on semantic segmentation and rely solely on RGB data, which captures color and geometric structure but lacks intrinsic material properties. This limitation can lead to issues like

*Project lead; † Work done during an internship at KAUST.

metamerism, where two objects appear identical under a particular light source despite having different materials and spectral radiance distributions. Furthermore, many materials exhibit distinct spectral behaviors outside the visible spectral range, such as in the near-infrared, where we can find indicators for vegetation health [21], or in the ultraviolet, where cues about minerals can be extracted from fluorescence [22]; information that RGB data fails to capture.

Spectral imaging (SI) emerges as a powerful framework for scene representation, as it provides rich spectral information across multiple wavelengths [5, 12, 15]. This, in turn, allows for a more complete characterization of a scene, effectively encoding unique spectral cues for different materials. Integrating spectral data into implicit neural representations remains a challenge. Standard NeRF formulations model scene appearance using only RGB values, which inherently limits their capacity to differentiate materials with similar color but distinct spectral properties. While recent efforts have attempted to extend NeRF to spectral data [3, 26] by replacing the standard RGB output with per-channel spectral radiance predictions, these approaches treat spectral information as an additional dimension in the output space without explicitly leveraging its inherent structure, e.g. materials have representative spectral signatures.

In real-world scenes, the spectral signature at a given point often arises from a mixture of multiple materials rather than a single homogeneous substance. Spectral unmixing [17] is the process of decomposing these mixed signals into their constituent materials, represented as a set of pure spectral signatures (*endmembers*), and estimating their corresponding *abundances*, which define the proportion of each endmember at a given 3D point. Consequently, spectral unmixing serves as a fundamental tool for distinguishing materials within a scene [44]. However, to the best of our knowledge, there is currently no work in view synthesis that fully exploits the rich structure of spectral data, not only to enhance the quality of novel view generation but also to enable more precise and physically meaningful identification of materials within a scene.

Contributions. In this paper, we propose UnMix-NeRF, the first framework that integrates spectral unmixing into NeRF, allowing joint spectral novel view synthesis and unsupervised material segmentation. Our method learns global endmembers through a dictionary optimized during training and models spectral reflectance via a diffuse and specular decomposition. The diffuse component is represented by per-point abundances and scaling factors, while a dedicated branch predicts view-dependent specular effects. The final spectral radiance is obtained by combining both components. Moreover, the per-point spectral unmixing naturally enables unsupervised material segmentation by leveraging the rendered abundance vectors. To summarize, our key contributions are:

- (i) We introduce UnMix-NeRF, a framework for hyperspectral novel view synthesis that leverages spectral unmixing to achieve high-quality spatial and spectral renderings by encoding material properties from a multi-view 3D hyperspectral dataset.
- (ii) Our framework inherently enables unsupervised segmentation by leveraging the rendered abundance vectors learned during training, which accurately delineate spatially varying material properties within scenes.
- (iii) Exploiting the learned endmember dictionary, our framework supports scene editing through direct and customizable manipulation of scene appearance with fine-grained control over individual materials.

Our approach not only outperforms state-of-the-art methods in hyperspectral view synthesis but also achieves pixel-accurate 3D material segmentation. Also, we extend a synthetic hyperspectral dataset with paired 3D material annotations, enabling rigorous evaluation of spectral unmixing and segmentation in complex scenes.

2. Related Work

Neural Radiance Fields. NeRF revolutionized novel view synthesis by representing scenes as continuous neural fields. Through volumetric rendering and multi-view optimization, it learns to predict density and color at any 3D point and viewing direction, enabling photorealistic view synthesis [31]. Several works have extended NeRF to new imaging modalities: SeaThru-NeRF [23] models underwater light scattering, Thermal-NeRF [45] handles infrared emissions, LiDAR-NeRF [39] enables LiDAR point cloud synthesis, and NeSpof [18] incorporates polarization information. For hyperspectral imaging, HS-NeRF [3] and Spec-NeRF [24] directly extend NeRF to predict per-wavelength radiance. Recent approaches like SpectralNeRF [26], introduce physically-based spectral rendering by extending NeRF with an additional autoencoder to predict per-wavelength radiance, and HyperGS [40], a concurrent work that adapts 3D Gaussian Splatting for spectral scene reconstruction. However, these methods treat spectral bands as independent channels, failing to take advantage of the rich material information encoded within the spectral signatures. These extensions require substantial computational overhead and also ignore the structured nature of spectral information, limiting their ability to disentangle materials.

Spectral Unmixing. Spectral unmixing decomposes pixels into endmembers (materials signatures) and their abundances (percentage of material per pixel) [17]. Beyond its use in remote sensing, unmixing enables material understanding in indoor and outdoor scenes in robotics [6], industrial inspection [9], and material classification [47]. The process typically involves identifying the different concurrent endmember spectra and their fractional abundances, per pixel, while enforcing physical constraints like non-

negativity and sum-to-one (i.e., the sum of all abundances must be 100%). Classical approaches employ the linear mixing model (LMM), which assumes a linear combination of endmembers [17]. However, LMM struggles with spectral variability caused by illumination, atmospheric conditions, and material mixing effects [2]. Extensions like the Augmented LMM (ALMM) [13] and the Extended LMM (ELMM) [42] address this by incorporating scaling factors and environmental variations through learned dictionaries and bundle-based representations, improving abundance estimation accuracy under real-world conditions. Recent deep learning approaches leverage transformers and convolutional neural networks (CNNs) to learn endmembers and abundances directly from data [10, 46]. While these methods have shown promise for 2D spectral unmixing, they are currently limited to single-view analysis and have not been extended to 3D reconstruction, where material decomposition could enable richer scene understanding. The challenge of incorporating physical unmixing constraints into neural architectures for 3D scenes remains unexplored.

Material Segmentation. Recent advances in self-supervised learning and foundation models have revolutionized unsupervised segmentation. Methods like DINO [33] and SAM [19] learn powerful visual features that allow zero-shot segmentation without task-specific training. These approaches leverage large-scale pretraining to discover semantic boundaries and object-centric representations, reducing the need for manual annotations. Materialistic [37] builds on this foundation, enabling interactive material selection based on a user-provided pixel, showing robustness to shading and geometric variations. For NeRF, several works explore unsupervised object segmentation through multi-view consistency. RFP [29] and ONeRF [28] segment objects by propagating visual features through the radiance field using photometric consistency. More recent approaches like LERF [16] and OpenNeRF [7] leverage language-vision models to enable open-vocabulary 3D segmentation. However, these methods focus on semantic segmentation rather than material understanding, as they rely on RGB appearance features that can be ambiguous for material discrimination. Recent works like MaterialSeg3D [27] attempt material segmentation in 3D but still operate on RGB images using semantic priors learned from labeled datasets. Concurrent work SAMa [8] requires a user-provided click and is limited to RGB data. Hyperspectral imaging, with its rich spectral information, naturally enables unsupervised material segmentation due to the inherent discriminative power of spectral signatures [34], thus eliminating the ambiguity associated with RGB features.

3. Preliminary

In this section, we review the NeRF volumetric rendering formulation and provide an overview of spectral unmixing.

3.1. Neural Radiance Fields

In NeRF, a scene is represented as a continuous volumetric function F parameterized by a Multi-Layer Perceptron (MLP). For any 3D point $\mathbf{x} = (x, y, z)$ and viewing direction $\mathbf{d} = (\theta, \phi)$, the network predicts a density σ and a view-dependent RGB color $\mathbf{c}$ as $(\sigma, \mathbf{c}) = F(\mathbf{x}, \mathbf{d})$. To render an image from a novel viewpoint, NeRF performs volumetric rendering by sampling points along each camera ray, defined as $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$, where $\mathbf{o} \in \mathbb{R}^3$ is the ray origin and $\mathbf{d} \in \mathbb{R}^3$ is the unit direction vector. The radiance is then accumulated using the classical volume rendering equation [31]. In practice, this integral is approximated as

$$C(\mathbf{r}) = \sum_{i=1}^{N_t} T_i (1 - \exp(-\sigma_i \delta_i)) \mathbf{c}_i, \quad (1)$$

where $T_i = \exp(-\sum_{j=1}^{i-1} \sigma_j \delta_j)$, $\delta_i = t_{i+1} - t_i$ is the distance between adjacent samples, and N_t is the total number of samples. The network is optimized by minimizing the mean squared error between rendered and ground truth pixel colors across multiple training views:

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

where $\mathcal{R}$ represents the set of camera rays from all training views and C^* denotes the corresponding ground truth.

3.2. Spectral Unmixing

A hyperspectral image with B spectral bands and $H \times W$ spatial pixels can be represented as $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_n, \dots, \mathbf{y}_N] \in \mathbb{R}^{B \times N}$, where $\mathbf{y}_n \in \mathbb{R}^{B \times 1}$ is a spectral pixel, for $n = 1, \dots, N$, with $N = HW$ as the total number of pixels. Let $\mathbf{E} = [\mathbf{e}_1, \dots, \mathbf{e}_k, \dots, \mathbf{e}_K] \in \mathbb{R}^{B \times K}$ be the endmember matrix, containing the spectral signatures of K pure materials. The Linear Mixing Model (LMM) [17] assumes that each spectral pixel is a linear combination of these endmembers,

$$\mathbf{y}_n = \mathbf{E} \mathbf{a}_n + \boldsymbol{\epsilon}_n, \quad (3)$$

where $\mathbf{a}_n \in \mathbb{R}^K$ is the abundance vector associated with pixel n -th, quantifying the contribution of each endmember and $\boldsymbol{\epsilon}_n$ is the residual (error approximation) term. The abundance coefficients satisfy the following constraints:

$$\mathbf{a}_n \succeq 0, \quad \mathbf{1}^T \mathbf{a}_n = 1, \quad (4)$$

ensuring that abundances are non-negative and that they sum to one, enforcing a physically meaningful representation of spectral mixing.

To account for spectral variability caused by the environment and illumination, the Extended Linear Mixing Model (ELMM) [42] introduces pixel-wise scaling factors $\mathbf{S}_n \in \mathbb{R}^{K \times K}$,

$$\mathbf{y}_n = \mathbf{E} \mathbf{S}_n \mathbf{a}_n + \boldsymbol{\epsilon}_n, \quad (5)$$

where $\mathbf{S}_n$ is a diagonal matrix with non-negative entries.




Figure 2. **Overview of NeRF unmixing approach.** Given a viewing direction (θ, ϕ) and 3D coordinates (x, y, z) , our network predicts a density σ and a set of abundance vectors $\tilde{\mathbf{a}}$, along with scalar factors $\tilde{\mathbf{s}}$ that scale the spectral endmembers $\mathbf{E}$. The scaled endmembers, when multiplied by the predicted abundances $\tilde{\mathbf{a}}$, yield the diffuse reflectance component $\tilde{\mathbf{c}}_d$. In parallel, a specular reflectance $\tilde{\mathbf{c}}_s$ and an additional tint factor h are also predicted to capture local illumination effects. The final spectral radiance is obtained by combining these diffuse and specular components $\tilde{\mathbf{c}}_d + \tilde{\mathbf{c}}_s$ under volumetric rendering with the predicted density. An ℓ_2 loss on the reconstructed spectral signatures is sufficient to guide unmixing and specular-diffuse modeling.

4. Proposed Method

Figure 2 illustrates our framework for hyperspectral novel view synthesis and material segmentation, which integrates spectral unmixing with volumetric rendering. In our approach, given a viewing direction (θ, ϕ) and a 3D point (x, y, z) , the network predicts a density σ together with a set of abundance vectors and scalar factors that scale a set of spectral endmembers. These scaled endmembers, when combined with the predicted abundances, form the diffuse reflectance component $\tilde{\mathbf{c}}_d$. In parallel, the model estimates a specular reflectance $\tilde{\mathbf{c}}_s$ and an additional tint factor h to account for local illumination effects. By blending these diffuse and specular components through volumetric rendering, the framework generates the final spectral radiance, which can be mapped from the scene coordinate space to sRGB for image synthesis (see, e.g., [11, 43]). Moreover, the per-point spectral unmixing naturally enables unsupervised material segmentation by leveraging the rendered abundance vectors.

4.1. Spectral Unmixing Field

We extend the standard NeRF formulation by incorporating the ELMM at each 3D point (see section 3.2). Let $\mathbf{x} = (x, y, z)$ be a coordinate in the scene. We first predict a density value along with an intermediate latent feature via an MLP, denoted as $(\sigma, \mathbf{z}) = F_\sigma(\mathbf{x})$, which maps $\mathbf{x}$ to a scalar density $\sigma \geq 0$ and a latent feature vector $\mathbf{z} \in \mathbb{R}^D$, where D is the dimensionality of the feature vector. This density is used for volumetric rendering (Section 4.3), while $\mathbf{z}$ serves as input for subsequent predictions.

Next, we introduce a learnable dictionary of K endmembers $\mathbf{E} \in \mathbb{R}^{B \times K}$, where each column $\mathbf{e}_k \in \mathbb{R}^B$ repre-

sents the pure spectral signature of the k -th material. These global endmembers are jointly optimized with the network parameters and can be initialized either by standard 2D unmixing methods (e.g., using vertex component analysis (VCA) [32]) or randomly.

To model local spectral mixing, we predict two components at each point: (i) a set of K *scaling factors* to account for illumination and environmental variability, and (ii) a set of K *abundances* representing the fractional presence of each material. Specifically, for (i), a dedicated MLP is used to obtain the scaling factors as $\mathbf{s} = F_s(\mathbf{x})$, with $\mathbf{s} \in \mathbb{R}^K$. To ensure non-negativity, we apply a sigmoid function to $\mathbf{s}$, hence obtaining $\tilde{\mathbf{s}} = \text{sigmoid}(\mathbf{s})$, with each element $\tilde{s}_k \in [0, 1]$. (ii) The intermediate feature $\mathbf{z}$ obtained from $F_\sigma(\cdot)$ is then used by an abundance head $F_a(\cdot)$ to predict raw abundances and a tint factor (see Section 4.2), expressed as $(\mathbf{a}, h) = F_a(\mathbf{x})$, where $\mathbf{a} \in \mathbb{R}^K$ and $h \in \mathbb{R}$. A softmax activation with temperature $\tau > 0$ is applied to $\mathbf{a}$, and computed as follows:

$$\tilde{a}_k = \frac{\exp(a_k/\tau)}{\sum_{j=1}^K \exp(a_j/\tau)}, \quad (6)$$

thus ensuring non-negativity and the sum-to-one constraint according to Eq. (4). Additionally, h is passed through a sigmoid activation to constrain its values within $[0, 1]$, thus modulating the contribution of the diffuse color. Each abundance $\tilde{a}_k$ corresponds to the fractional presence of the k -th material at point $\mathbf{x}$, providing a physically meaningful decomposition of the current spectral signature. With these, the diffuse reflectance is computed as follows:

$$\tilde{\mathbf{c}}_d = \mathbf{E} \tilde{\mathbf{S}} \tilde{\mathbf{a}}, \quad (7)$$

where $\tilde{\mathbf{S}}$ is a diagonal matrix with non-negative scaling factors $\tilde{s}$ on its diagonal.

4.2. Specular Field

Since spectral unmixing is inherently designed for diffuse reflectance, it fails to capture view-dependent effects such as specular highlights. To address this limitation, we introduce a specular field. Motivated by Ref-NeRF [43], we leverage the tint factor h predicted by the abundance head $F_a(\cdot)$ (see Section 4.1) to modulate the specular reflectance, allowing for better control of view-dependent effects. Then, taking both $\mathbf{x}$ and $\mathbf{d}$ as input, an MLP $F_c(\cdot)$ is used to obtain the specular reflectance vector as $\mathbf{c}_s = F_c(\mathbf{x}, \mathbf{d})$, where $\mathbf{c}_s \in \mathbb{R}^B$. The tint factor is then used to modulate this specular term as

$$\tilde{\mathbf{c}}_s = h\mathbf{c}_s. \quad (8)$$

The final per-point spectral reflectance is obtained by combining the diffuse and specular components under the dichromatic model [4]:

$$\mathbf{c} = \mathbf{c}_d + \tilde{\mathbf{c}}_s. \quad (9)$$

4.3. Volume Rendering and Camera Response

Following the standard NeRF formulation (see Section 3.1), we sample points along camera rays and accumulate the spectral radiance. The accumulated spectral radiance along a ray is computed using Eq. (1). Similarly, we can volume-render the material abundances to obtain the per-ray abundance vectors as:

$$\mathbf{A}(\mathbf{r}) = \sum_{i=1}^{N_t} T_i (1 - \exp(-\sigma_i \delta_i)) \tilde{\mathbf{a}}_i, \quad (10)$$

for the i -th sample location $\mathbf{x}_i$ along the ray $\mathbf{r}$. This allows us to render the abundances for each point in the scene

To recover an RGB image from the rendered hyperspectral signature, we employ a spectral response model. This conversion involves transforming the spectral power distribution in the scene coordinate space to a device-independent color space (such as XYZ) and then to the target sRGB color space. Let $\mathbf{M} \in \mathbb{R}^{3 \times B}$ be the camera spectral response matrix that encapsulates both the color matching functions and the display primaries; then, the RGB color is given by

$$\mathbf{C}_{\text{rgb}}(\mathbf{r}) = \mathbf{M} \mathbf{C}(\mathbf{r}). \quad (11)$$

This approach ensures accurate color reproduction by modeling the entire imaging pipeline from scene radiance to displayed color.

4.4. Material Segmentation via Cluster Probe

For unsupervised material segmentation, we perform a cluster assignment for each ray's predicted spectral signature

by employing the learned spectral endmembers $\mathbf{E}$ as cluster centers. Given the predicted spectral signature for each ray $\mathbf{C}(\mathbf{r}) \in \mathbb{R}^B$, we compute the normalized inner product against the dictionary of endmembers as

$$\mathbf{p}(\mathbf{r}) = \text{softmax} \left(\frac{\mathbf{E}^\top \mathbf{C}(\mathbf{r})}{\|\mathbf{E}\| \|\mathbf{C}(\mathbf{r})\|} \right), \quad (12)$$

where $\mathbf{p}(\mathbf{r}) \in \mathbb{R}^K$ contains the probabilities indicating cluster memberships for each of the K materials. Finally, material segmentation is obtained by assigning each ray $\mathbf{r}$ to the material cluster with the highest probability following

$$m(\mathbf{r}) = \arg \max_k p_k(\mathbf{r}). \quad (13)$$

4.5. Loss Function

We jointly optimize all parameters, including the endmember matrix $\mathbf{E}$, and the weights of the MLPs for density (F_σ), scaling (F_s), abundance (F_a), and specular reflectance (F_c), by minimizing a combination of hyperspectral and RGB reconstruction errors. Let $\mathbf{C}(\mathbf{r})$ and $\mathbf{C}_{\text{rgb}}(\mathbf{r})$ denote the rendered hyperspectral and RGB predictions for ray $\mathbf{r}$, and let $\mathbf{C}^*(\mathbf{r})$ and $\mathbf{C}_{\text{rgb}}^*(\mathbf{r})$ be the corresponding ground-truth measurements. The losses are defined as

$$L_{\text{spec}} = \sum_{\mathbf{r} \in \mathcal{R}} \|\mathbf{C}(\mathbf{r}) - \mathbf{C}^*(\mathbf{r})\|_2^2, \quad (14)$$

$$L_{\text{rgb}} = \sum_{\mathbf{r} \in \mathcal{R}} \|\mathbf{C}_{\text{rgb}}(\mathbf{r}) - \mathbf{C}_{\text{rgb}}^*(\mathbf{r})\|_2^2, \quad (15)$$

where $\mathcal{R}$ is the set of all training rays. Then, our final objective is given by

$$L = \lambda_{\text{spec}} L_{\text{spec}} + \lambda_{\text{rgb}} L_{\text{rgb}}, \quad (16)$$

with hyperparameters λ_{spec} and λ_{rgb} balancing the importance of hyperspectral versus RGB fidelity. Minimizing this loss not only leads to accurate novel view synthesis but also ensures physically consistent spectral unmixing. This results in both high-quality spectral reconstructions and reliable material segmentation through the per-ray abundance predictions.

4.6. Implementation Details

We implemented our method within the Nerfstudio framework [38], building on top of the Nerfacto implementation. The hidden feature dimension D was set to 16, and we integrated Nerfacc [25] into the volumetric rendering and sampling process to reduce training time. Additionally, we applied gradient scaling based on the squared ray distance to each pixel, as suggested in [35]. For optimization, we used the Adam optimizer with a learning rate of 1×10^{-2} and an exponential learning rate scheduler. All models were

trained for 20,000 iterations, and all experiments were performed on an NVIDIA A100 GPU, with a maximum memory usage of approximately 20 GB. For all experiments, we set $\lambda_{\text{spec}} = 5$ and $\lambda_{\text{rgb}} = 1$. To ensure physically meaningful spectral signatures, the endmembers $\mathbf{E}$ were constrained at each iteration by applying element-wise clamping within the range $[0, 1]$. For each trained scene, the number of endmembers K is selected as a heuristic approximation of the number of distinct materials present in the scene.

5. Experimental Results

We extensively evaluate the performance of our proposed approach using standard metrics and compare it with state-of-the-art methods.

5.1. Datasets, Preprocessing, Metrics and Baselines

Datasets. Three different datasets are employed for hyperspectral novel view synthesis and material segmentation:

- **NeSpoF Dataset [18].** We use the NeSpoF spectro-polarimetric dataset, focusing exclusively on the S_0 Stokes component representing spectral radiance. Specifically, we utilize the synthetic dataset comprising 21 spectral bands ranging from 450 nm to 650 nm. This dataset includes four distinct scenes: *ajar*, *hotdog*, *cbox_dragon*, and *cbox_sphere*. We adopt the provided camera poses and their calibrated RGB projections. Additionally, we extended the dataset by generating material maps from the synthetic scenes, obtaining ground-truth labels for material segmentation. Further details about the extension are provided in the supplementary material.
- **Surface Optics Dataset [3].** The Surface Optics dataset contains real-world imagery captured across 128 spectral bands, spanning wavelengths from 370 nm to 1100 nm. This dataset consists of four object-centric scenes: Rosemary, Basil, Tools and Origami. We adopt the original camera poses.
- **BaySpec Dataset [3].** The BaySpec dataset offers hyperspectral data across 140 spectral bands, covering a wavelength range from 400 nm to 1110 nm. BaySpec offers three scenes: Pinecone, Caladium and Anacampseros and includes measured camera poses as described in the original dataset publication.

Metrics. We evaluate UnMix-NeRF and competing methods on hyperspectral view synthesis using rigorous metrics that capture both accuracy and spectral fidelity, including PSNR, SSIM, Spectral Angle Mapping (SAM) [20], and RMSE. For unsupervised material segmentation, we report standard metrics, including mean Intersection over Union (mIoU) and F1 score.

Baselines. We compare our method against existing novel view methods available for hyperspectral imaging. Specifically, we compare with NeSpoF [18], originally designed

Ablation Step	PSNR $\uparrow$	SSIM $\uparrow$	SAM $\downarrow$	RMSE $\downarrow$
Base: Unmix Only	24.73	0.710	0.063	0.058
+ Physical Constraint	22.36	0.722	0.075	0.076
+ Scaling Factors	29.36	0.914	0.026	0.034
+ RGB Loss	30.92	0.923	0.025	0.028
+ VCA Initialization	31.09	0.923	0.024	0.027
+ Specular Field	33.20	0.935	0.022	0.023

Table 1. Ablation study on the NeSpoF dataset, evaluating the impact of each component in UnMix-NeRF. Metrics include PSNR ($\uparrow$), SSIM ($\uparrow$), Spectral Angle Mapper (SAM, $\downarrow$), and RMSE ($\downarrow$). Each row sequentially adds a component to the baseline model. The final configuration, incorporating all elements, achieves the best overall performance.

for spectro-polarimetric imaging, restricting our comparison to the spectral component only (S_0 Stokes). Also, we include HS-NeRF, which extends the NeRF framework from the conventional RGB 3-channel output to a multi-channel spectral representation. Lastly, we also compare against HyperGS, a framework for hyperspectral novel view synthesis based on the 3D Gaussian Splatting (3DGS). Also, we include comparisons with traditional NeRF models adapted for hyperspectral data (NeRF, MipNeRF, TensorRF, Nerfacto and MipNeRF-360) for the Bayspec and Surface Optics datasets.

5.2. Ablation Studies

To evaluate the effectiveness and impact of individual components in our UnMix-NeRF model, we performed extensive ablation studies. Table 1 summarizes the results obtained from each ablation step performed over the NeSpoF dataset. The initial baseline model (Unmix Only) predicts abundances using global endmembers without further constraints or enhancements. We then sequentially added:

1. Physical constraints, from Section 3.1, using activation functions to enforce realistic abundance estimates.
2. Scaling factors, in line with the ELMM, enable the model to better represent local spectral variability.
3. An RGB loss term consisting of a spectral-to-RGB projection that helps with realism and spectral consistency.
4. The VCA initialization of the global endmembers instead of random initialization.
5. Finally, the specular field component is used to accurately model specular reflectances via a dichromatic model.

Results in Table 1 illustrate the incremental contributions of each component added to our UnMix-NeRF baseline. The initial model (Unmix Only) achieves moderate spectral reconstruction quality (PSNR = 24.73), establishing a performance baseline without additional constraints. Introducing physical constraints reduces the PSNR slightly by a margin of -2.37 due to the restrictive nature of these constraints on spectral variability. The addition of scaling fac-




Figure 3. Visualization of the top-4 performing methods for frame 51 out of 359 for the Caladium plant scene from the Bayspec dataset. The top row shows the 70th spectral channel out of the 141-channel predicted image, and the bottom row provides a raw pixel-wise mean relative absolute error heatmap of the scene.

tors significantly boosts the performance (PSNR increases by +7.00, SAM decreases by -0.049 , and RMSE decreases by -0.042), underscoring their importance in modeling local spectral variations. Incorporating the RGB loss further improves the spectral accuracy (PSNR improves by +1.56), demonstrating the benefit of coupling hyperspectral predictions with RGB supervision. Using the VCA initialization provides a minor improvement of +0.17 in PSNR, indicating that an informed initialization of the endmembers aids the optimization stability. Finally, integrating the specular field yields the best overall results (PSNR increases by +2.11, SSIM reaches 0.935), confirming that explicitly modeling specular reflectance substantially enhances spectral rendering and overall image quality.

5.3. Qualitative Results

To further assess the effectiveness of our UnMix-NeRF framework, we provide qualitative comparisons against state-of-the-art methods in hyperspectral novel view synthesis. Figure 3 illustrates the quality of reconstruction between different methods in the Caladium scene of the Bayspec dataset. The top row displays the 70th spectral channel from the reconstructed hyperspectral cube, while the bottom row depicts the corresponding mean relative absolute error (MRAE) heatmaps, highlighting pixel-wise reconstruction errors. Our method achieves the most accurate spectral predictions, significantly reducing reconstruction artifacts and preserving fine-grained spectral details.

Additionally, Figure 1 demonstrates the capabilities of UnMix-NeRF for material segmentation and scene editing. Leveraging the learned endmembers and their corresponding abundance maps, our approach enables precise material separation and targeted modifications to scene appearance. Furthermore, Figure 4 presents a visualization of the learned material abundances for the *Anacampseros* scene in the BaySpec dataset. Each image corresponds to a learned

Method	Scene	PSNR $\uparrow$	RMSE $\downarrow$	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.




Figure 4. Visualization of the learned material abundance maps for the *Anacampseros* scene from the BaySpec dataset (top row) and the *Ajar* scene from the NeSpoF dataset (bottom row). Each map represents the spatial distribution of a material’s abundance, highlighting distinct materials in the scene.

abundance map, highlighting the presence of distinct materials in the scene. The spatial variations observed across the different maps indicate that our approach successfully captures and disentangles material-specific contributions.

Bayspec Dataset

Method	Pinecone				Caladium				Anacampseros			
	PSNR $\uparrow$	SSIM $\uparrow$	SAM $\downarrow$	RMSE $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	SAM $\downarrow$	RMSE $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	SAM $\downarrow$	RMSE $\downarrow$
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

Table 3. Quantitative results on the BaySpec dataset for the *Pinecone*, *Caladium*, and *Anacampseros* scenes. Our method achieves the best performance across all metrics.

Method	Rosemary				Basil			
	PSNR $\uparrow$	SSIM $\uparrow$	SAM $\downarrow$	RMSE $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	SAM $\downarrow$	RMSE $\downarrow$
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

Table 4. Quantitative results on the Surface Optics dataset for the *Rosemary* and *Basil* scenes. Our method achieves the best overall performance across all metrics. Results for the *Tools* and *Origami* scenes are provided in the supplementary material.

5.4. Quantitative Results

Quantitative evaluations of our UnMix-NeRF method on three different datasets demonstrate superior performance across a variety of scenes and conditions. In all tables, the best performing result for each metric is highlighted in *red*, the second-best in *orange*, and the third-best in *yellow*.

Table 2 presents the quantitative comparison on the NeSpoF dataset, where our method achieves state-of-the-art performance while maintaining significantly lower computational cost. Our approach reaches an average PSNR of 33.2 while reducing the computation time to only 44 minutes. Additionally, we match the RMSE of NeSpoF, which requires over 11.9 hours per scene. The per-scene results highlight the robustness of our method.

We further evaluate UnMix-NeRF on the BaySpec dataset, as shown in Table 3. Our method consistently outperforms all baselines in all scenes, achieving the highest PSNR, SSIM, and the lowest SAM and RMSE values. Notably, our model improves over HyperGS by +3.15 PSNR on the *Caladium* scene. Similar trends are observed in the *Pinecone* and *Anacampseros* scenes, where our approach provides both higher reconstruction fidelity and better spectral consistency. These results show that explicitly modeling spectral unmixing in NeRF improves spectral prediction accuracy over per-pixel regression.




Figure 5. Unsupervised material segmentation of the *Ajar* scene.

Table 4 reports the results on the Surface Optics dataset. Our method achieves the best performance in both the *Rosemary* and *Basil* scenes, obtaining a +2.14 and +3.91 PSNR improvement over HyperGS, respectively. Furthermore, our approach achieves the lowest SAM and RMSE values, confirming its ability to accurately reconstruct spectral details across diverse material compositions. The results for the remaining scenes are provided in the supplementary material.

Additionally, we evaluate the effectiveness of our method for unsupervised material segmentation on the extended NeSpoF synthetic dataset. Our approach achieves an average F1 score of 0.41 and a mIoU of 0.28 across all scenes in the dataset, demonstrating UnMix-NeRF’s ability to extract meaningful material clusters without requiring supervision. Figure 5 shows the unsupervised material segmentation results of our method on the *Ajar* scene.

6. Conclusions

We introduced UnMix-NeRF, a novel approach integrating spectral unmixing into neural radiance fields to achieve simultaneous hyperspectral novel view synthesis and unsupervised material segmentation. By modeling spectral reflectance through learned global endmembers and per-point abundances, our method effectively captures intrinsic material properties. Our extensive evaluations show that UnMix-NeRF outperforms existing methods in spectral reconstruction quality and unsupervised material segmentation accuracy. Additionally, our framework supports intuitive material-based scene editing through direct manipulation of endmember dictionaries.

Acknowledgments. The research reported in this publication was supported by funding from King Abdullah University of Science and Technology (KAUST) - Center of Excellence for Generative AI, under award number 5940.

References

- [1] S Adarsh, S Mohamed Kaleemuddin, Dinesh Bose, and KI Ramachandran. Performance comparison of infrared and ultrasonic sensors for obstacles of different materials in vehicle/robot navigation applications. In *IOP Conference Series: Materials Science and Engineering*, page 012141. IOP publishing, 2016. 1
- [2] Ricardo Augusto Borsoi, Tales Imbiriba, José Carlos Moreira Bermudez, Cédric Richard, Jocelyn Chanussot, Lucas Drumetz, Jean-Yves Tournet, Alina Zare, and Christian Jutten. Spectral variability in hyperspectral data unmixing: A comprehensive review. *IEEE geoscience and remote sensing magazine*, 9(4):223–270, 2021. 3
- [3] Gerry Chen, Sunil Kumar Narayanan, Thomas Gautier Ottou, Benjamin Missaoui, Harsh Muriki, Cédric Pradalier, and Yongsheng Chen. Hyperspectral neural radiance fields. *arXiv preprint arXiv:2403.14839*, 2024. 2, 6
- [4] Robert L Cook and Kenneth E. Torrance. A reflectance model for computer graphics. *ACM Transactions on Graphics (ToG)*, 1(1):7–24, 1982. 5
- [5] Claudia V Correa, Carlos A Hinojosa, Gonzalo R Arce, and Henry Arguello. Multiple snapshot colored compressive spectral imager. *Optical Engineering*, 56(4):041309–041309, 2017. 2
- [6] Sibren Dieters. Robot-aided hyperspectral imaging for mineral exploration in underground mining environments. 2024. 2
- [7] Francis Engelmann, Fabian Manhardt, Michael Niemeyer, Keisuke Tateno, Marc Pollefeys, and Federico Tombari. Opennerf: open set 3d neural scene segmentation with pixel-wise features and rendered novel views. *arXiv preprint arXiv:2404.03650*, 2024. 1, 3
- [8] Michael Fischer, Iliyan Georgiev, Thibault Groueix, Vladimir G Kim, Tobias Ritschel, and Valentin Deschaintre. Sama: Material-aware 3d selection and segmentation. *arXiv preprint arXiv:2411.19322*, 2024. 3
- [9] Nahum Gat, Suresh Subramanian, Jacob Barhen, and Nikzad Toomarian. Spectral imaging applications: remote sensing, environmental monitoring, medicine, military operations, factory automation, and manufacturing. In *25th AIPR Workshop: Emerging Applications of Computer Vision*, pages 63–77. SPIE, 1997. 2
- [10] Preetam Ghosh, Swalpa Kumar Roy, Bikram Koirala, Behnood Rasti, and Paul Scheunders. Hyperspectral unmixing using transformer network. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–16, 2022. 3
- [11] Peter Hedman, Pratul P Srinivasan, Ben Mildenhall, Jonathan T Barron, and Paul Debevec. Baking neural radiance fields for real-time view synthesis. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5875–5884, 2021. 4
- [12] Carlos Hinojosa, Esteban Vera, and Henry Arguello. A fast and accurate similarity-constrained subspace clustering algorithm for hyperspectral image. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 14:10773–10783, 2021. 2
- [13] Danfeng Hong, Naoto Yokoya, Jocelyn Chanussot, and Xiao Xiang Zhu. An augmented linear mixing model to address spectral variability for hyperspectral unmixing. *IEEE Transactions on Image Processing*, 28(4):1923–1938, 2018. 3
- [14] Benran Hu, Junkai Huang, Yichen Liu, Yu-Wing Tai, and Chi-Keung Tang. Nerf-rpn: A general framework for object detection in nerfs. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 23528–23538, 2023. 1
- [15] Longqian Huang, Ruichen Luo, Xu Liu, and Xiang Hao. Spectral imaging with deep learning. *Light: Science & Applications*, 11(1):61, 2022. 2
- [16] Justin Kerr, Chung Min Kim, Ken Goldberg, Angjoo Kanazawa, and Matthew Tancik. Lrf: Language embedded radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19729–19739, 2023. 1, 3
- [17] Nirmal Keshava and John F Mustard. Spectral unmixing. *IEEE signal processing magazine*, 19(1):44–57, 2002. 2, 3
- [18] Youngchan Kim, Wonjoon Jin, Sunghyun Cho, and Seung-Hwan Baek. Neural spectro-polarimetric fields. In *SIGGRAPH Asia 2023 Conference Papers*, pages 1–11, 2023. 2, 6
- [19] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4015–4026, 2023. 1, 3
- [20] Fred A Kruse, Adam B Lefkoff, Joseph W Boardman, Kathleen B Heidebrecht, AT Shapiro, PJ Barloon, and Alexander FH Goetz. The spectral image processing system (sips)—interactive visualization and analysis of imaging spectrometer data. *Remote sensing of environment*, 44(2-3): 145–163, 1993. 6
- [21] Neha Kureel, Jyoti Sarup, Shafique Matin, Suresh Goswami, and Kapil Kureel. Modelling vegetation health and stress using hypersepctral remote sensing data. *Modeling Earth Systems and Environment*, pages 1–16, 2022. 2
- [22] Joseph R Lakowicz. *Principles of fluorescence spectroscopy*. Springer, 2006. 2
- [23] Deborah Levy, Amit Peleg, Naama Pearl, Dan Rosenbaum, Derya Akkaynak, Simon Korman, and Tali Treibitz. Seathru-nerf: Neural radiance fields in scattering media. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 56–65, 2023. 2
- [24] Jiabao Li, Yuqi Li, Ciliang Sun, Chong Wang, and Jinhui Xiang. Spec-nerf: Multi-spectral neural radiance fields. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2485–2489. IEEE, 2024. 2

- [25] Ruilong Li, Hang Gao, Matthew Tancik, and Angjoo Kanazawa. Nerfacc: Efficient sampling accelerates nerfs. *arXiv preprint arXiv:2305.04966*, 2023. 5
- [26] Ru Li, Jia Liu, Guanghui Liu, Shengping Zhang, Bing Zeng, and Shuaicheng Liu. Spectralnerf: Physically based spectral rendering with neural radiance field. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 3154–3162, 2024. 2
- [27] Zeyu Li, Ruitong Gan, Chuanchen Luo, Yuxi Wang, Jiaheng Liu, Ziwei Zhu, Qing Li, Xucheng Yin, Man Zhang, Zhaoxiang Zhang, et al. Materialseg3d: Segmenting dense materials from 2d priors for 3d assets. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 370–379, 2024. 3
- [28] Shengnan Liang, Yichen Liu, Shangzhe Wu, Yu-Wing Tai, and Chi-Keung Tang. Onerf: Unsupervised 3d object segmentation from multiple views. *arXiv preprint arXiv:2211.12038*, 2022. 3
- [29] Xinhang Liu, Jiaben Chen, Huai Yu, Yu-Wing Tai, and Chi-Keung Tang. Unsupervised multi-view object segmentation using radiance field propagation. *Advances in Neural Information Processing Systems*, 35:17730–17743, 2022. 3
- [30] Yichen Liu, Benran Hu, Chi-Keung Tang, and Yu-Wing Tai. Sanerf-hq: Segment anything for nerf in high quality. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3216–3226, 2024. 1
- [31] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. 1, 2, 3
- [32] José MP Nascimento and José MB Dias. Vertex component analysis: A fast algorithm to unmix hyperspectral data. *IEEE transactions on Geoscience and Remote Sensing*, 43(4):898–910, 2005. 4
- [33] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 3
- [34] Fabian Perez and Hoover Rueda-Chacón. Beyond appearances: Material segmentation with embedded spectral information from rgb-d imagery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 293–301, 2024. 3
- [35] Julien Philip and Valentin Deschaintre. Floaters No More: Radiance Field Gradient Scaling for Improved Near-Camera Training. In *Eurographics Symposium on Rendering*. The Eurographics Association, 2023. 5
- [36] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 1
- [37] Prafull Sharma, Julien Philip, Michaël Gharbi, Bill Freeman, Fredo Durand, and Valentin Deschaintre. Materialistic: Selecting similar materials in images. *ACM Transactions on Graphics*, 42(4), 2023. 3
- [38] Matthew Tancik, Ethan Weber, Evonne Ng, Ruilong Li, Brent Yi, Terrance Wang, Alexander Kristoffersen, Jake Austin, Kamyar Salahi, Abhik Ahuja, et al. Nerfstudio: A modular framework for neural radiance field development. In *ACM SIGGRAPH 2023 conference proceedings*, pages 1–12, 2023. 5
- [39] Tang Tao, Longfei Gao, Guangrun Wang, Yixing Lao, Peng Chen, Hengshuang Zhao, Dayang Hao, Xiaodan Liang, Mathieu Salzmann, and Kaicheng Yu. Lidar-nerf: Novel lidar view synthesis via neural radiance fields. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 390–398, 2024. 2
- [40] Christopher Thirgood, Oscar Mendez, Erin Chao Ling, Jon Storey, and Simon Hadfield. Hypergs: Hyperspectral 3d gaussian splatting. *arXiv preprint arXiv:2412.12849*, 2024. 2
- [41] Giuseppe Vecchio, Renato Sortino, Simone Palazzo, and Concetto Spampinato. Matfuse: controllable material generation with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4429–4438, 2024. 1
- [42] Miguel Angel Veganzones, Lucas Drumetz, Guillaume Tochon, Mauro Dalla Mura, Antonio Plaza, J Bioucas-Dias, and Jocelyn Chanussot. A new extended linear mixing model to address spectral variability. In *2014 6th Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing (WHISPERS)*, pages 1–4. IEEE, 2014. 3
- [43] Dor Verbin, Peter Hedman, Ben Mildenhall, Todd Zickler, Jonathan T Barron, and Pratul P Srinivasan. Ref-nerf: Structured view-dependent appearance for neural radiance fields. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5481–5490. IEEE, 2022. 4, 5
- [44] Erin B Wetherley, Dar A Roberts, and Joseph P McFadden. Mapping spectrally similar urban materials at sub-pixel scales. *Remote Sensing of Environment*, 195:170–183, 2017. 2
- [45] Tianxiang Ye, Qi Wu, Junyuan Deng, Guoqing Liu, Liu Liu, Songpengcheng Xia, Liang Pang, Wenxian Yu, and Ling Pei. Thermal-nerf: Neural radiance fields from an infrared camera. *arXiv preprint arXiv:2403.10340*, 2024. 2
- [46] Xiangrong Zhang, Yujia Sun, Jingyan Zhang, Peng Wu, and Licheng Jiao. Hyperspectral unmixing via deep convolutional neural networks. *IEEE Geoscience and Remote Sensing Letters*, 15(11):1755–1759, 2018. 3
- [47] Min Zhao, Jie Chen, and Zhe He. A laboratory-created dataset with ground truth for hyperspectral unmixing evaluation. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 12(7):2170–2183, 2019. 2
- [48] Junwei Zheng, Jiaming Zhang, Kailun Yang, Kunyu Peng, and Rainer Stiefelhagen. Materobot: Material recognition in wearable robotics for people with visual impairments. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2303–2309. IEEE, 2024. 1