LINUU: LINEARIZATION- AND INVERSION-BASED NON-LINEAR UNSUPERVISED UNMIXING

Xander Haijen, Xuanwen Tao, Bikram Koirala, Paul Scheunders

Imec-Visionlab, Department of Physics, University of Antwerp, Belgium

ABSTRACT

Nonlinear spectral mixing poses a significant challenge in hyperspectral unmixing, as standard linear assumptions often fail in real-life scenarios. While the Multilinear Mixing Model (MLM) effectively models higher-order photon interactions through a pixel-specific nonlinearity parameter, existing unsupervised approaches struggle with inaccurate endmember initializations and unreliable parameter estimation. In this paper, we propose **LINUU** (Linearization- and Inversion-based Non-linear Unsupervised Unmixing), a novel physics-inspired framework for unsupervised nonlinear unmixing. LINUU leverages the Linearizing Space Transform to map the observed nonlinearly mixed hyperspectral image into a linearized space where low-rank matrix factorization holds. To tackle the non-convex optimization of the nonlinearity parameter, an interior-point method with an L-BFGS Hessian approximation is employed. Experimental results on synthetic and real hyperspectral datasets demonstrate that LINUU outperforms state-of-the-art methods in accurately estimating endmembers, abundances, and nonlinearity parameters.

***Index Terms*—** Hyperspectral unmixing, nonlinear mixing, multilinear mixing model, linearizing space transform, nonnegative low rank factorization

1. INTRODUCTION

In spectral unmixing, the goal is to decompose mixed hyperspectral pixels, which contain several distinct materials, into constituent materials, which are called Endmembers (EMs), and their fractional abundances. The most prevalent model used in spectral unmixing is the Linear Mixing Model (LMM), which models a pixel as a linear combination of its constituent spectra [1]. However, *nonlinear* spectral mixing effects are often a crucial component in many real-world remote sensing scenarios. These include the presence of intimate mixtures, vegetation canopies and urban scenes, where light may undergo several material interactions before reaching the sensor. To handle these complex scenarios, several nonlinear mixing models have been introduced. One of these is the Multilinear Mixing Model (MLM) [2], which models a possibly infinite number of interaction events by introducing a pixel-wise interaction probability parameter.

With these augmented and improved models, it is possible to perform unmixing in scenarios that do not conform to the assumptions of the LMM. However, most of these methods are supervised, meaning they require precise EM estimates in order to obtain reliable abundances. If these are not readily available, they must be estimated from the image. This can be obtained a priori (using geometric endmember extraction algorithms) or using joint optimization techniques such as nonnegative matrix factorization (NMF). However, the geometric interpretation that underpins these methods does not hold in nonlinear cases [1], which renders the EM estimates unreliable and consequently prevents precise abundance estimation. In the context of blind MLM unmixing, several methods have been proposed, but they fail to address the challenges mentioned above, or their architectures lack physical explainability, which is crucial in physics-inspired unmixing.

In this paper, we propose **LINUU: Linearization- and Inversion-based Non-linear Unsupervised Unmixing**, an unsupervised unmixing framework that jointly estimates the abundances, the EMs and the nonlinearity parameters. It utilizes the Linearizing Space Transform (LST) proposed in [3] and the nonnegative low-rank property of linear unmixing to construct a low-rank linearized image, which is then unmixed using a standard linear unmixing strategy.

2. RELATED WORK

Let P denote the number of spectral bands, N the number of pixels in an image, and K the number of EMs. We then use $\mathbf{y}_n$ to denote the n-th pixel, $\mathbf{e}_k$ for the k-th EM and a_{kn} is the fractional abundance of the k-th EM in the n-th pixel. We will use the corresponding boldcase and uppercase notation for matrices, e.g., $\mathbf{Y}$ is the image matrix.

2.1. Linear Mixing Model

The LMM is the most fundamental model in spectral unmixing and is given by the equation

$$\mathbf{y}_n = \mathbf{E}\mathbf{a}_n + \boldsymbol{\epsilon}_n \tag{1}$$

The work of Bikram Koirala was supported by the Research Foundation Flanders, Belgium (FWO: 1250824N-7028). Xander Haijen is a PhD fellow of the Research Foundation Flanders, Belgium (FWO: 1103226N).

where $\boldsymbol{\epsilon}_n \in \mathbb{R}^P$ is additive noise. The abundances must satisfy the physical Abundance Nonnegativity Constraint (ANC) and Abundance Sum-to-one Constraint (ASC). For an entire image, the formula becomes

$$\mathbf{Y} = \mathbf{E}\mathbf{A} + \mathbf{N}. \tag{2}$$

where $\mathbf{N}$ is a noise matrix. This means that $\mathbf{Y}$ is approximately a low-rank matrix and can be factorized into two nonnegative factor matrices $\mathbf{A}$ and $\mathbf{E}$. This *nonnegative low-rank* property is a fundamental and crucial algebraic property of the linear mixing equation.

2.2. Multilinear Mixing Model (MLM)

The MLM uses the ray-based approximation of light and explicitly models higher-order interactions in a pixel. It assumes that the ray of incident light interacts with a pure material, after which it has a probability of p of undergoing further interactions and a probability of $1-p$ to escape the scene and reach the sensor [2]. This leads to an infinite series representation, which can be re-written as

$$\mathbf{y}_n = \frac{(1-p_n)\mathbf{E}\mathbf{a}_n}{1-p_n\mathbf{E}\mathbf{a}_n} + \boldsymbol{\epsilon}_n. \tag{3}$$

p_n is a pixel-specific nonlinearity coefficient. A large p_n indicates strong nonlinearity and many higher-order interactions, and $p_n = 0$ indicates purely linear mixing. Note that p_n can also be negative. In this case, the probability interpretation no longer holds, but Eq. (3) nevertheless remains a valid mixing equation. In the original MLM paper [2], the authors proposed using the loss function

$$\left\| \mathbf{y}_n - \frac{(1-p_n)\mathbf{E}\mathbf{a}_n}{1-p_n\mathbf{E}\mathbf{a}_n} \right\|_2^2 \tag{4}$$

and optimize it w.r.t. p_n and $\mathbf{a}_n$ for every pixel, while enforcing $p_n < 1$ and while $\mathbf{a}_n$ is subject to the ASC and ANC. This method succeeds in accurately unmixing nonlinear scenes, but only when the provided EMs are highly accurate.

Later, the authors in [4] proposed an unsupervised unmixing method based on the MLM, NLU-MLM, which uses a gradient projection method. The authors in [5] proposed a similarity graph-based method for MLM unmixing, G-MLM. These methods, however, struggle when the initialization of the EMs is inaccurate.

Recently, a renewed interest in MLM-based unmixing using deep learning (DL) has emerged, alongside a more general interest in DL-based nonlinear unmixing. In [6], the authors propose a multiscale dilated transformer-based architecture for jointly estimating all parameters. In [3], the authors propose MPGU-Net, a physics-guided network that consists of a three-block architecture, with one block for estimating the abundances, one for estimating the EMs, and one for estimating the nonlinearity parameters. Despite their improvements over early MLM unmixing methods, the DL methods still fail to accurately estimate the nonlinearity parameters, and they lack clear physical interpretations due to their black-box architectures.

2.3. Linearizing Space Transform

The authors in [3] use the minimal volume constraint – a popular technique in linear unmixing – and adapt it for use in the nonlinear case with MPGU-Net. This is then used in the loss function for network training. This process uses the so-called Linearizing Space Transform (LST). It starts with a Taylor series expansion of the nonlinear model, which is then analytically inverted to find the linear component of the model in terms of the nonlinear spectrum and the nonlinearity parameter. Let $\mathbf{x}_n = \mathbf{E}\mathbf{a}_n$ be the linear component of the model, then the expansion for the MLM is given as

$$\begin{aligned}
\mathbf{y}_n &= (1-p_n)\mathbf{x}_n + (1-p_n)p_n\mathbf{x}_n \odot \mathbf{x}_n \\
&\quad + (1-p_n)p_n^2\mathbf{x}_n \odot \mathbf{x}_n \odot \mathbf{x}_n + \dots \\
&= (1-p_n)\mathbf{x}_n + p_n\mathbf{x}_n \odot \big((1-p_n)(\mathbf{x}_n + p_n\mathbf{x}_n \odot \mathbf{x}_n \\
&\quad + p_n^2\mathbf{x}_n \odot \mathbf{x}_n \odot \mathbf{x}_n + \dots\big) \\
&= (1-p_n)\mathbf{x}_n + p_n\mathbf{x}_n \odot \mathbf{y}_n,
\end{aligned} \tag{5}$$

where $\odot$ represents elementwise multiplication. This can be re-written as

$$\mathbf{x}_n = \frac{\mathbf{y}_n}{1-p_n+p_n\mathbf{y}_n} \tag{6}$$

where all operations are performed elementwise. This means that, when the nonlinearity parameter p_n is known, it is possible to linearize the image matrix, and then perform standard linear unmixing on this new image to obtain the abundances.

3. PROPOSED METHOD

In this section, we described the proposed LINUU method in detail. LINUU jointly estimates a low nonnegative rank factorization of the image and the nonlinearity parameter via a BCD scheme. Although this paper focuses on MLM unmixing, the framework can easily be adapted for unmixing with other nonlinear mixing models that admit an LST formulation (see [3] for examples).

3.1. Model formulation

For simplicity of notation, we introduce the following definitions for the forward MLM and the inverse MLM

$$\mathrm{MLM}(\mathbf{x}_n, p_n) = \frac{(1-p_n)\mathbf{x}_n}{1-p_n\mathbf{x}_n} \tag{7}$$

$$\mathrm{MLM}^{-1}(\mathbf{y}_n, p_n) = \frac{\mathbf{y}_n}{1-p_n+p_n\mathbf{y}_n}, \tag{8}$$

where all operators are elementwise. LINUU minimizes the difference between the LST-transformed input data and a low-rank matrix $\widehat{\mathbf{X}} = \mathbf{WH}$:

$$\begin{aligned} \min_{\mathbf{p},\mathbf{W},\mathbf{H}} \ & \frac{1}{2} \left\| \mathrm{MLM}^{-1}(\mathbf{Y}, \mathbf{p}) - \mathbf{WH} \right\|_F^2 \\ \text{s.t. } & \mathbf{W} \geq 0, \quad \mathbf{H} \geq 0, \quad \mathbf{H}^\top \mathbf{1} = \mathbf{1} \\ & p_{\min} \leq p_n < 1, \ n = 1, 2, \ldots, N \end{aligned} \tag{9}$$

where $\mathbf{p} \in \mathbb{R}^N$ is the vector of all p_n. $\mathbf{W} \in \mathbb{R}^{P \times K}$ and $\mathbf{H} \in \mathbb{R}^{K \times N}$ are the factor matrices, which force the expression $\mathrm{MLM}^{-1}(\mathbf{Y}, \mathbf{p})$ to conform to a matrix of nonnegative rank at most K, aligning with the algebraic interpretation of linear mixing. $\mathrm{MLM}^{-1}(\mathbf{Y}, \mathbf{p})$ is the column-wise application of MLM^{-1}:

$$\begin{aligned} &\mathrm{MLM}^{-1}(\mathbf{Y}, \mathbf{p}) = \\ &\left[\mathrm{MLM}^{-1}(\mathbf{y}_1, p_1) \ \mathrm{MLM}^{-1}(\mathbf{y}_2, p_2) \ \cdots \ \mathrm{MLM}^{-1}(\mathbf{y}_N, p_N)\right] \end{aligned} \tag{10}$$

It is important to note that, even though $\mathbf{W}$ and $\mathbf{H}$ are structurally similar to the EM and abundance matrices and have similar constraints, $\mathbf{W}$ and $\mathbf{H}$ are not the final abundance and EM estimates. They are merely the mathematical representation of the low nonnegative rank structure that we wish to impose on $\mathrm{MLM}^{-1}(\mathbf{Y}, \mathbf{p})$.

After an optimum for Eq. (9) is found, a sequential approach is used for finding the EMs and abundances. First, we use a geometric EM extraction algorithm, namely SISAL [7], on the linearized image matrix $\widehat{\mathbf{X}} = \mathbf{WH}$. This leads to the final EM estimate. Using this estimate, we then perform unmixing using Fully Constrained Least Squares Unmixing (FCLSU), leading to the final abundance estimate.

3.2. Optimization Procedure

For minimizing Eq. (9), we use a BCD framework that alternately updates $\mathbf{H}$, $\mathbf{W}$ and $\mathbf{p}$.

3.2.1. $\mathbf{H}$ *and* $\mathbf{W}$ *updates*

For updating $\mathbf{H}$, we note that the problem is convex w.r.t. $\mathbf{H}$ and enjoys an analytical solution. By setting the derivative to zero and solving for $\mathbf{H}$, we get the update rule

$$\mathbf{H} \leftarrow (\mathbf{W}^\top \mathbf{W})^{-1} (\mathbf{W}^\top \mathbf{Z}) \tag{11}$$

where $\mathbf{Z} = \mathrm{MLM}^{-1}(\mathbf{Y}, \mathbf{p})$. Lastly, we project $\mathbf{H}$ onto the simplex using

$$\mathbf{h}_n \leftarrow \frac{\max(\mathbf{h}_n, 0)}{\| \max(\mathbf{h}_n, 0) \|_1}, \ n = 1, 2, \ldots, N \tag{12}$$

For updating $\mathbf{W}$, we use an analogous approach to $\mathbf{H}$. The analytical update rule is given by

$$\mathbf{W} \leftarrow (\mathbf{Z}\mathbf{H}^\top)(\mathbf{H}\mathbf{H}^\top)^{-1} \tag{13}$$

and constraint projection is performed by $\mathbf{W} \leftarrow \max(\mathbf{W}, 0)$.

3.2.2. $\mathbf{p}$ *update*

Unlike the other two updates, the update for the parameter $\mathbf{p}$ does not enjoy an analytical solution. In fact, due to the highly nonlinear nature of the inverse model $\mathrm{MLM}^{-1}(\mathbf{Y}, \mathbf{p})$, this is a highly irregular and nonconvex problem. Standard first-order methods do not suffice for solving this problem, so instead we use second order methods. They incorporate second order derivative (Hessian) information which greatly increases performance [8]. We use an interior-point (IP) algorithm and approximate the Hessian by using the L-BFGS approximation. The IP algorithm naturally limits the values of $\mathbf{p}$ to the allowed range. Since the gradient of $\mathrm{MLM}^{-1}(\mathbf{Y}, \mathbf{p})$ can be computed analytically, both the memory and computational cost of the IP solver remain acceptable.

The parameter $\mathbf{p}$ is prone to overfitting: when it has reached a good value that aligns the inverse model closely to a low-rank matrix, it will start to adapt to either model noise present in $\mathbf{Y}$, or it will start fluctuating to accommodate the changes in $\mathbf{W}$ and $\mathbf{H}$ as they change throughout iterations. This creates a limit cycle where all variables are constantly changing to match the target provided by the other variables, which are also changing. This is exacerbated by the highly nonconvex inverse model. To prevent this overfitting we fix p_n once it reached an acceptable value. In every iteration, we check the forward loss for each pixel

$$\ell_n = \| \mathbf{y}_n - \mathrm{MLM}(\hat{\mathbf{x}}_n, p_n) \|_2^2 \tag{14}$$

where $\hat{\mathbf{x}}_n$ is the n-th column of $\widehat{\mathbf{X}}$. When the value for ℓ_n does not improve for a user-defined number of iterations $T_\ell > 0$, the value for p_n is frozen and not updated again.

3.2.3. Convergence

Lastly, we perform a convergence check in every iteration. When the relative changes in the constituent variables all fall below a user-defined threshold $\epsilon > 0$, the algorithm terminates. Note that, after termination, the EMs and abundances are determined by using a linear unmixing strategy applied to the linearized image matrix $\widehat{\mathbf{X}} = \mathbf{WH}$.

4. EXPERIMENTS

We compare our newly proposed LINUU method to several nonlinear unmixing methods using the MLM: (1) **MLM**: the original MLM algorithm [2] with EMs provided by SISAL, (2) **G-MLM**: a graph-based method by [5] with EMs provided by SISAL, (3) **NLU-MLM**: an unsupervised BCD-based algorithm [4], (4) **DTU-Net**: a transformer-based DL architecture [6], and (5) **MPGU-Net**: a three-block network incorporating the LST into the loss function [3]. For DTU-Net and MPGU-Net, implementations were available from the respective papers. For the other methods, we use our own implementations that closely follow the descriptions in

Table 1. Performance comparison of various unmixing methods on synthetic GRF data. The best results are highlighted in bold.

Method	EM SAD (°)	A RMSE	p RMSE	Time (s)
MLM	46.788	0.2561	0.2596	**1.90**
G-MLM	14.520	0.2991	59.618	116
NLU-MLM	6.0378	0.1075	0.6302	50.2
DTU-Net	9.2948	0.0630	0.2734	16.1
MPGU-Net	5.4004	0.0615	0.5940	13.2
LINUU	**1.2575**	**0.0288**	**0.0520**	26.1

the respective papers. For the EMs, we use the Spectral Angle Distance (SAD) in degrees as the error metric, and for the abundances and nonlinearity parameter, we use the Root Mean Square Error (RMSE). For LINUU, we set $p_{\min} = -1$ in both experiments.

4.1. Synthetic data

In the first experiment, we use four endmembers from the USGS Spectral Library [9], which have a spectral range from 350 nm to 2,500 nm. We select 224 equidistant bands for further use. We generate the abundance maps using Gaussian Random Fields (GRFs) that comply with the ANC and ASC. We also generate a nonlinearity parameter field using GRFs with values limited to the interval $[-1, 1)$. Using these synthetic abundances and nonlinearity parameters, we generate a synthetic image conforming to the MLM and add Gaussian noise with an SNR of 30 dB. We use SISAL to generate the initial EM guesses (including the initial guess for $\mathbf{W}$ in LINUU). The total image consists of 100×100 pixels.

Table 1 shows the performance by the different methods, alongside their runtime. MLM and G-MLM perform poorly, since they rely heavily on the SISAL EMs, which are highly inaccurate for nonlinear cases. The DL methods and LINUU have a better EM and abundance estimation than older methods, and have an acceptable runtime. Of these methods, LINUU has the best EM and abundance errors, and is the only method that generates an acceptable estimate for the nonlinearity map, with an RMSE of about 5%, with all other methods having an error of 25% or higher. Fig. 1 shows the estimated nonlinearity parameter maps for the different methods. DTU-Net produces $\mathbf{p}$ estimates which are very small, so it assumes near-linear mixing throughout the image. MPGU-Net produces a uniform guess of 0.5, rendering the nonlinearity parameter map completely uninformative. Other methods also struggle with finding the right scale for $\mathbf{p}$, and only LINUU finds the right magnitude and is able to recover the structure of the nonlinearity parameter map accurately.

4.2. Real image

In this second experiment, we perform unmixing on a subset of the Houston dataset. This image comprises 144 spectral bands with a spectral range from 380 nm to 1,050 nm. The image contains a large stadium, alongside a tree-lined entrance. These are two types of structures that often cause light to interact with several materials before reaching the sensor, so a degree of nonlinearity can be expected in this image. Four ground truth EMs, which were manually extracted from the image, are provided with the dataset. We again use SISAL to extract the initial EM estimates. Table 2 shows the EM estimation error for the different methods. LINUU performs best overall, with a low SAD for all EMs. The abundance maps are given in Fig. 2. LINUU clearly identifies the red roofs located at the corners of the stadium and along the entrance, and also identifies the vegetation both inside and outside of the stadium. Several methods struggle to unmix the

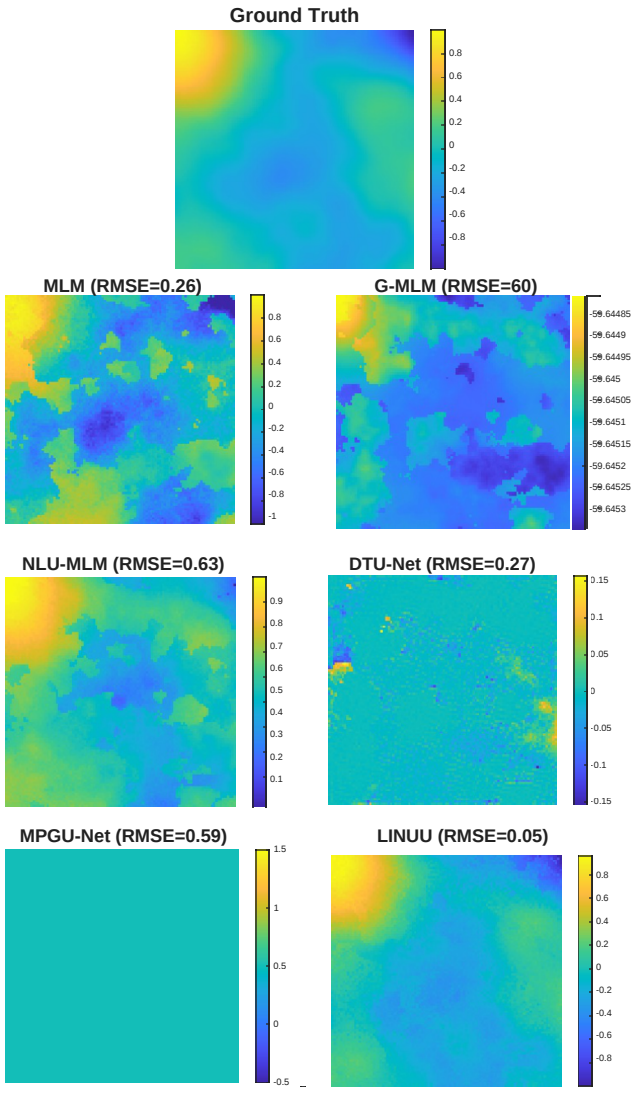


Fig. 1. Nonlinearity parameter maps for different methods and the ground truth. Note the scale of the colorbars, especially for DTU-Net and G-MLM.

EM	MLM	G-MLM	NLU-MLM	DTU-Net	MPGU-Net	LINUU
Grass	12.564	10.272	7.1852	5.3875	4.2092	**3.221**
Red Roofs	9.2551	7.8211	10.399	23.039	6.3944	**3.8185**
Concrete	3.8712	14.263	**3.3971**	11.665	27.712	5.6467
Asphalt	38.517	**5.8783**	10.517	29.923	10.698	9.2514
Mean	16.052	9.5585	7.8746	17.503	12.253	**5.4844**

Table 2. EM SADs (in degrees) for different methods. The best results are highlighted in bold.

image, e.g. G-MLM and MPGU-Net, which produce uniform abundance maps, and the same abundance maps four times, respectively. This demonstrates that LINUU is able to unmix complex images and provide physically meaningful abundance maps, alongside accurate EM estimates.

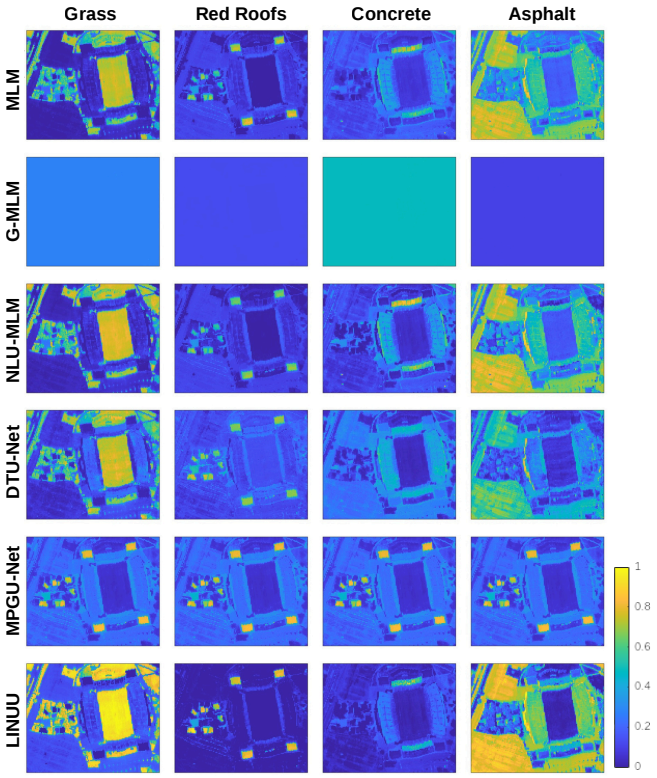


Fig. 2. Abundance maps obtained by unmixing the Houston image for different methods.

5. CONCLUSION

In this paper, we proposed LINUU, a novel physics-inspired framework for unsupervised nonlinear unmixing. It addresses the current lack of unsupervised nonlinear unmixing methods that can effectively deal with imprecise EM estimates, while simultaneously providing good estimates for the nonlinearity parameter. In this paper, we focused on MLM-based unmixing and showed that LINUU outperforms the state of the art. The LINUU framework is very flexible and can be expanded to other nonlinear mixing models as well. This is, however, left to future work.

6. REFERENCES

[1] José M. Bioucas-Dias et al., "Hyperspectral Unmixing Overview: Geometrical, Statistical, and Sparse Regression-Based Approaches," *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.*, vol. 5, no. 2, pp. 354–379, Apr. 2012.

[2] Rob Heylen and Paul Scheunders, "A Multilinear Mixing Model for Nonlinear Spectral Unmixing," *IEEE Trans. Geosci. Remote Sens.*, vol. 54, no. 1, pp. 240–251, Jan. 2016.

[3] Xuanwen Tao et al., "MPGU-Net: Multimodel Physics-Guided Network for Nonlinear Hyperspectral Unmixing," *IEEE Trans. Geosci. Remote Sens.*, vol. 64, pp. 5518522–5518522, 2026.

[4] Qi Wei et al., "Unsupervised Nonlinear Spectral Unmixing Based on a Multilinear Mixing Model," *IEEE Trans. Geosci. Remote Sens.*, vol. 55, no. 8, pp. 4534–4544, Aug. 2017.

[5] Minglei Li et al., "A Graph Regularized Multilinear Mixing Model for Nonlinear Hyperspectral Unmixing," *Remote Sens.*, vol. 11, no. 19, pp. 2188, Jan. 2019.

[6] Chentong Wang et al., "DTU-Net: A Multi-Scale Dilated Transformer Network for Nonlinear Hyperspectral Unmixing," *IEEE Trans. Geosci. Remote Sens.*, vol. 64, pp. 1–17, 2026.

[7] Jose M. Bioucas-Dias, "A variable splitting augmented Lagrangian approach to linear spectral unmixing," in *2009 First Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing*, Aug. 2009.

[8] Jorge Nocedal and Stephen J. Wright, *Numerical Optimization*, Springer Series in Operations Research and Financial Engineering. Springer New York, 2006.

[9] Raymond F. Kokaly et al., "USGS Spectral Library Version 7," Tech. Rep. 1035, U.S. Geological Survey, 2017.