Physics-Guided Regime Unmixing

Paula Pacheco[1,2,3], Pablo Granitto[4], Juan B. Cabral[1,2,3]
paulapacheco@mi.unc.edu.ar

[1] GVT-CONAE, Centro Espacial Teófilo Tabanera, Argentina
[2] CONICET, Argentina
[3] FAMAF-UNC, Córdoba, Argentina
[4] CIFASIS, CONICET–UNR, Rosario, Argentina

Abstract. The Linear Mixing Model (LMM) dominates spectral unmixing for its simplicity, but fails under multiple scattering; existing nonlinear models compensate by applying a fixed regime uniformly across entire scenes. We propose *Physics-Guided Regime Unmixing* (PGRU), which estimates a pixel-wise scalar $\xi_i \in [0, 1]$ from observable physical features to activate nonlinear mixing only where justified. Residuals from the Generalized Bilinear Model (GBM), the Post-Nonlinear Mixing Model (PPNM), and Hapke are combined via learned attention, yielding interpretable regime maps. Experiments on Samson, Jasper Ridge, and Urban show consistent improvements over baselines, with physical coherence $\rho > 0.90$.

Keywords: hyperspectral unmixing, nonlinear mixing, physics-guided learning, regime selection, remote sensing

Régimen de Desmezclado Guiado por Física

Abstract. El Modelo de Mezcla Lineal (LMM) predomina en el desmezclado espectral por su simplicidad, pero falla ante la dispersión múltiple; los modelos no lineales existentes compensan esto aplicando un régimen fijo de manera uniforme en toda la escena. Proponemos el *Physics-Guided Regime Unmixing* (PGRU), que estima un escalar por píxel $\xi_i \in [0, 1]$ a partir de características físicas observables para activar la mezcla no lineal únicamente donde está justificado. Los residuos del Modelo Bilineal Generalizado (GBM), el Modelo de Mezcla Post-No Lineal (PPNM) y Hapke se combinan mediante atención aprendida, generando mapas de régimen interpretables. Los experimentos en Samson, Jasper Ridge y Urban muestran mejoras consistentes frente a las líneas base, con coherencia física $\rho > 0.90$.

Keywords: desmezclado hiperespectral, mezcla no lineal, aprendizaje guiado por física, selección de régimen, teledetección

1 Introduction

Hyperspectral sensors capture surface reflectance across hundreds of contiguous spectral bands, enabling fine-grained material characterization for applications including agriculture, geology, and environmental monitoring (Bhargava et al., 2024). Because spatial resolution is typically coarser than the scale of surface heterogeneity, each pixel contains the mixed spectral response of multiple materials. Spectral unmixing recovers the pure signatures (endmembers) and their fractional contributions (abundances) from this mixture.

The LMM (Keshava & Mustard, 2002) is the standard approach due to its simplicity and physical interpretability, but assumes photons interact with a single material before reaching the sensor. This assumption breaks down under multiple scattering, motivating nonlinear extensions such as the Generalized Bilinear Model (GBM) (Halimi et al., 2011), the Post-Nonlinear Mixing Model (PPNM) (Altmann et al., 2012), and the Hapke radiative transfer model (Hapke, 1981). These models improve reconstruction accuracy in strongly nonlinear scenes, but apply nonlinearity uniformly across all pixels, without distinguishing which pixels actually require it or why.

Deep learning approaches, primarily autoencoders and convolutional networks (Palsson et al., 2018, 2022), achieve strong reconstruction performance, but operate as black boxes and offer no physical explanation for regime selection.

We address this gap with PGRU, which introduces a pixel-wise regime parameter $\xi_i \in [0, 1]$ guided by observable physical features. Unlike previous approaches, PGRU does not require labels indicating whether pixels are linear or nonlinear: regime assignment emerges from the tension between reconstruction improvement and physical feature consistency.

2 Physics-Guided Regime Unmixing

Each pixel $i \in \{1, \ldots, N\}$ is reconstructed as $\hat{y}_i = S_{\text{lin},i} + \xi_i \, \delta_{\text{nl},i}$, where $S_{\text{lin},i} = \sum_{m=1}^{M} a_{im} e_m$ is the standard LMM component, with e_m the m-th endmember spectrum, $a_{im} \geq 0$ the abundance coefficients satisfying $\sum_m a_{im} = 1$, and $\delta_{\text{nl},i}$ a nonlinear residual weighted by $\xi_i \in [0, 1]$. When $\xi_i \approx 0$ the model reduces to LMM; when $\xi_i \approx 1$ the nonlinear contribution is fully active.

The residual $\delta_{\text{nl},i} = \sum_k \alpha_{ik} \, \delta_{ik}$ combines three physical models via learned attention weights α_{ik} obtained through a temperature-controlled softmax: GBM captures pairwise scattering between endmembers, PPNM introduces quadratic distortion over the linear mixture, and Hapke describes multiple scattering from radiative transfer theory. Entropy regularization on α_{ik} encourages concentration toward the dominant physical mechanism at each pixel.

The central contribution is that ξ_i is guided by observable scene properties rather than driven solely by reconstruction error.

$$\xi_i = \sigma\left(\sum_k w_k F_{ik} + b\right) \tag{1}$$

where i indexes pixels, k indexes features, F_{ik} are physical and geometric features of pixel i (spectral curvature, NDVI, NDVI gradient, EMP, DMP, LBP), w_k are learned weights, b is a learned bias, and σ is the sigmoid function. This is equivalent to a logistic regression over the feature space: the sign and magnitude of w_k directly reveal which scene properties predict nonlinear behavior and in which direction, making the regime decision for any pixel a readable sum of feature contributions.

To train the model, a per-pixel reconstruction gain is defined as $\Delta_{\mathrm{res},i} = \|y_i - S_{\mathrm{lin},i}\| - \|y_i - \hat{y}_i\|$, and the full objective is:

$$\mathcal{L} = \sum_i \Big[-\tanh(\Delta_{\mathrm{res},i}) + \lambda_{\mathrm{feat}}(\xi_i - \xi_{\mathrm{prior},i})^2 + \lambda_{\mathrm{sp}}\, \xi_i \nabla^2 \xi_i \Big] + \lambda_w \|w\|^2 \qquad (2)$$

where $\lambda_{\mathrm{feat}}, \lambda_{\mathrm{sp}}, \lambda_w \geq 0$ are regularization coefficients, and $\xi_{\mathrm{prior},i}$ is the feature-based prior. The first term rewards nonlinear activation where it reduces reconstruction error, saturated by tanh for large gains. The second anchors ξ_i to the feature-based prior, with λ_{feat} annealed during training. The third promotes spatial coherence via a Laplacian regularizer. No pixel labels are used: regime assignment emerges from reconstruction evidence and physical feature consistency.

Abundances a_{im} are estimated via projected gradient descent onto the unit simplex, enforcing non-negativity and the sum-to-one constraint, shared across all mixing models for fair comparison. The final regime map is calibrated via Platt scaling to align ξ_i with the empirical probability of nonlinear improvement.

3 Experiments

We evaluate PGRU on three widely used benchmark datasets (Zhu, 2017): **Samson** (95×95 pixels, 156 bands, 3 endmembers: water, vegetation, soil), **Jasper Ridge** (100×100, 198 bands, 4 endmembers: water, trees, soil, road), and **Urban** (307×307, 162 bands, 5 endmembers: asphalt, grass, trees, roof, soil). Reference endmembers provided with each dataset are used in all methods. Baselines are LMM (linear), GBM, and PPNM; all methods share the same abundance estimation procedure. Reconstruction is evaluated with Spectral Angle Distance (SAD), RMSE, and relative RMSE (rRMSE).

Table 1 reports reconstruction metrics across all datasets and methods. Preliminary experiments indicate that PGRU achieves lower reconstruction errors on all three datasets. The improvement is especially pronounced in Samson (rRMSE drops from 0.495 to 0.052) and Jasper Ridge (rRMSE from 0.347 to 0.068), with substantial gains also observed in the more complex Urban scene. Notably, the nonlinear baselines (GBM, PPNM) do not consistently outperform LMM, confirming that applying nonlinearity uniformly can be counterproductive. PGRU avoids this by activating nonlinearity only where justified.

Beyond reconstruction accuracy, we evaluate whether the regime activation maps are physically grounded. Physical coherence $\rho = \mathrm{corr}(\xi, \Delta_{\mathrm{res}})$ measures

Table 1. Reconstruction metrics on three benchmark datasets. Best results in **bold**.

Dataset	Metric	LMM	GBM	PPNM	PGRU
Samson	SAD	0.158	0.141	0.143	**0.052**
	RMSE	0.121	0.141	0.132	**0.013**
	rRMSE	0.495	0.578	0.540	**0.052**
Jasper Ridge	SAD	0.288	0.290	0.287	**0.151**
	RMSE	0.055	0.075	0.070	**0.011**
	rRMSE	0.347	0.476	0.444	**0.068**
Urban	SAD	0.121	0.120	0.115	**0.070**
	RMSE	0.067	0.087	0.081	**0.013**
	rRMSE	0.331	0.430	0.400	**0.064**

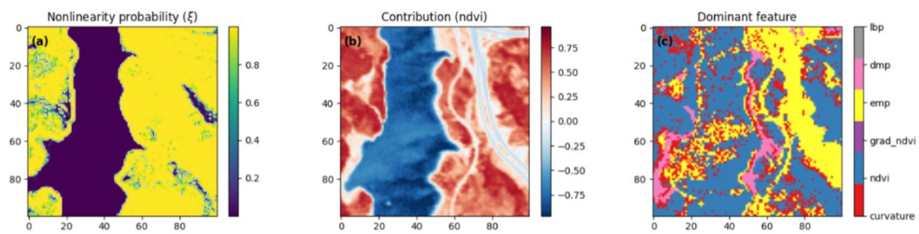

Fig. 1. Jasper Ridge: (a) nonlinearity map ξ, (b) NDVI contribution, (c) dominant feature per pixel. Feature-guided regime selection concentrates nonlinearity where multiple scattering is physically expected.

alignment between the regime map and per-pixel reconstruction gain, exceeding 0.90 across all datasets ($\rho = 0.93$, 0.91, 0.92 for Samson, Jasper Ridge, and Urban respectively).

The feature weight vector w reveals which physical properties predict nonlinear behavior in each scene. In Jasper Ridge, NDVI dominates: nonlinear activation is concentrated in dense vegetation, where multiple scattering is expected, and suppressed over the water body. In Urban, spectral curvature drives regime selection, reflecting structural complexity at material boundaries rather than vegetation-related interactions. This scene-adaptive behavior emerges without supervision.

Beyond global feature importance, PGRU provides pixel-level explainability through spatial maps that decompose regime selection into feature contributions. Fig. 1 illustrates this behavior for Jasper Ridge. The spatial organization of these explanations suggests that the model does not arbitrarily assign regimes, but anchors them to measurable scene properties, providing direct answers to the question *why is this pixel nonlinear?*. The large improvements in Table 1 reflect this selectivity: unlike GBM and PPNM, PGRU avoids the reconstruction penalty that uniform nonlinear activation incurs in linear regions.

4 Conclusion

We presented PGRU, which addresses a question often left open by prior unmixing methods: *why is this pixel nonlinear?* By grounding the regime scalar ξ_i in observable physical features, the model produces interpretable regime maps without labeled data. Initial results on three benchmarks show that selective, explainable nonlinearity outperforms both uniform linear and uniform nonlinear approaches. A current limitation is the lack of ablation: the relative contributions of the physical model ensemble, the feature-guided ξ_i, and spatial regularization remain to be disentangled. Future work will extend PGRU to blind unmixing, incorporate spatial context into regime estimation, and compare against recent deep learning baselines.

Acknowledgments. This work was partially supported by CONICET (Argentina). P.P. was supported by a fellowship from CONICET and CONAE.
Language editing assistance was provided by AI tools. All scientific content remains the responsibility of the authors.

References

Altmann, Y., Halimi, A., Dobigeon, N., & Tourneret, J.-Y. (2012). Supervised nonlinear spectral unmixing using a postnonlinear mixing model for hyperspectral imagery. *IEEE Transactions on Image Processing*, *21*(6), 3017–3025.

Bhargava, A., Sachdeva, A., Sharma, K., Alsharif, M. H., Uthansakul, P., & Uthansakul, M. (2024). Hyperspectral imaging and its applications: A review. *Heliyon*, *10*(12).

Halimi, A., Altmann, Y., Dobigeon, N., & Tourneret, J.-Y. (2011). Nonlinear unmixing of hyperspectral images using a generalized bilinear model. *IEEE Transactions on Geoscience and Remote Sensing*, *49*(11), 4153–4162.

Hapke, B. (1981). Bidirectional reflectance spectroscopy: 1. theory. *Journal of Geophysical Research: Solid Earth*, *86*(B4), 3039–3054.

Keshava, N., & Mustard, J. F. (2002). Spectral unmixing. *IEEE signal processing magazine*, *19*(1), 44–57.

Palsson, B., Sigurdsson, J., Sveinsson, J. R., & Ulfarsson, M. O. (2018). Hyperspectral unmixing using a neural network autoencoder. *IEEE Access*, *6*, 25646–25656.

Palsson, B., Sveinsson, J. R., & Ulfarsson, M. O. (2022). Blind hyperspectral unmixing using autoencoders: A critical comparison. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, *15*, 1340–1372.

Zhu, F. (2017). Hyperspectral unmixing: Ground truth labeling, datasets, benchmark performances and survey. *arXiv preprint arXiv:1708.05125.*