

Joint Endmember Representation and Abundance Mapping for Hyperspectral Image Analysis

Terry D. Peters

School of Electrical Engineering and Computer Science, Oregon State University, Corvallis,
OR, USA.

terry976@oregonstate.edu

Gaurav Bhatia

Department of Electrical Engineering and Computer Science, University of Kansas, Lawrence,
KS, USA.

gaurav.bhatia@ku.edu

Abstract

Hyperspectral image analysis relies on the accurate decomposition of mixed pixels into constituent spectral endmembers and their corresponding fractional abundances. Traditional approaches treat endmember extraction and abundance estimation as sequential, often independent, tasks, which can lead to information loss and suboptimal reconstruction. This paper proposes a systemic perspective on joint endmember representation and abundance mapping, arguing that a unified learning framework can more effectively capture the intrinsic spectral-spatial structure of hyperspectral data. We examine architectural trade-offs between end-to-end neural architectures and hybrid physics-informed models, emphasizing the role of representation bottlenecks, reconstruction fidelity, and spatial regularization. Infrastructure considerations for deploying such models on satellite and airborne platforms are discussed, including computational constraints, energy efficiency, and real-time inference requirements. Robustness to noise, spectral variability, and missing bands is analyzed through the lens of adversarial training and attention mechanisms. Fairness and policy implications arise when abundance maps inform land-use classification, resource allocation, or environmental monitoring; we highlight risks of bias propagation and the need for transparent governance frameworks. The paper further explores sustainability challenges related to large-scale training data requirements and model carbon footprint. By integrating insights from signal processing, machine learning, and socio-technical systems, we advocate for joint representation frameworks that are not only accurate but also deployable, equitable, and resilient. A case study on weak-signal unmixing illustrates the importance of handling rare materials and fine-scale abundance reconstruction. The conclusion outlines future research directions toward standardized benchmarks, interpretable architectures, and participatory model validation.

Keywords

Hyperspectral unmixing, endmember representation, abundance mapping, joint learning, socio-technical infrastructure, robustness, policy implications.

1. Introduction

Hyperspectral imaging captures reflectance or radiance across hundreds of narrow contiguous spectral bands, enabling detailed material identification in remote sensing, agriculture, mineralogy, and defense. A fundamental challenge in hyperspectral analysis is that the spatial

resolution of sensors often yields mixed pixels, wherein the measured spectrum is a composite of multiple pure materials (endmembers) weighted by their fractional coverage (abundances). Unmixing, the process of recovering endmember spectra and abundance maps, is therefore critical for downstream applications. Historically, unmixing has been performed as a two-stage pipeline: endmember extraction via geometrical algorithms such as vertex component analysis or N-FINDR, followed by abundance inversion using constrained least-squares or sparse regression [1], [2]. While effective in controlled settings, this sequential approach suffers from several limitations: errors in endmember selection propagate to abundance estimation, the model cannot exploit spatial context across pixels, and the representation capacity is bounded by linear mixing assumptions [3].

Recent advances in deep learning have spurred interest in joint endmember representation and abundance mapping, where a single network simultaneously learns spectral bases and pixel-wise mixing coefficients. This paradigm shift promises tighter integration of spectral and spatial information, better generalization to nonlinear mixtures, and potential for end-to-end optimization using reconstruction loss or variational objectives [4], [5]. However, the transition from sequential to joint frameworks introduces new architectural design choices, governance challenges, and deployment constraints that merit careful examination from a systems perspective. This paper provides a comprehensive analysis of joint representation learning for hyperspectral unmixing, focusing not only on algorithmic performance but also on the broader infrastructure, robustness, fairness, and policy dimensions that determine real-world viability.

The structure of the paper is as follows. Section 2 reviews related work on both traditional and learning-based unmixing, highlighting the evolution toward joint models. Section 3 examines architectural paradigms, including encoder-decoder networks, autoencoders, and attention mechanisms, and discusses trade-offs in representational capacity, computational cost, and physical plausibility. Section 4 analyzes robustness across noise, spectral variability, and incomplete observations, with reference to adversarial training and weak-signal detection [6]. Section 5 addresses deployment and sustainability, covering on-device inference, energy budgets, and data lifecycle management. Section 6 explores fairness, accountability, and policy implications when unmixing outputs inform critical decisions in environmental justice, resource allocation, and surveillance. The required reference [18] is discussed in the context of weak-signal representation learning for rare material detection. Section 7 concludes with synthetic insights and a research agenda.

2. Background and Related Work

Traditional unmixing algorithms rely on the linear mixing model, which assumes each mixed pixel spectrum is a convex combination of endmember spectra plus additive noise. Endmember extraction methods such as pixel purity index and N-FINDR search for extreme points in the spectral simplex [7]. Once endmembers are identified, fully constrained least-squares solves for nonnegative abundances that sum to one. These approaches are computationally efficient and interpretable but fail under nonlinear mixing scenarios caused by intimate mixtures, multiple scattering, or topographic effects [8]. Furthermore, they treat each pixel independently, disregarding spatial correlations.

Sparse regression-based unmixing relaxes the need for explicit endmember extraction by assuming a large spectral library and enforcing sparsity in abundance coefficients [1]. While this addresses some limitations, libraries may be incomplete or mismatch scene-specific conditions. Deep learning methods emerged as an alternative, with autoencoders learning a

low-dimensional representation that implicitly encodes endmembers. For example, a fully connected autoencoder with nonnegative constraints on the decoder weights can be interpreted as a nonlinear generalization of linear unmixing [9]. However, these models often lack spatial regularization and may produce physically implausible endmember spectra.

Joint endmember representation and abundance mapping refers to architectures that simultaneously learn a spectral dictionary and pixel-level mixing coefficients in a single optimization loop. Convolutional neural networks (CNNs) and transformers have been employed to incorporate spatial context, enabling the network to exploit neighborhoods for more robust abundance estimation [10]. Variational autoencoders introduce probabilistic formulations that capture uncertainty in both endmembers and abundances [11]. The work in [18] proposes a weak-signal representation learning framework with gated abundance reconstruction, specifically targeting scenarios where target materials occupy only a small fraction of pixels; this architecture uses state-space modeling and weak-signal attention fusion to enhance sensitivity to rare endmembers, illustrating the importance of specialized design for imbalanced mixture regimes.

Despite these advances, many joint models are evaluated on synthetic or well-controlled datasets, and their transferability to real sensor data with varying noise, atmospheric effects, and calibration artifacts remains an open challenge. Moreover, the governance of model deployment—particularly in defense and environmental monitoring—requires careful consideration of failure modes and interpretability. The subsequent sections analyze these systemic dimensions in depth.

3. Architectural Paradigms and Structural Trade-offs

Joint endmember and abundance learning must balance representational flexibility against physical interpretability. At one extreme, purely data-driven encoder-decoder architectures minimize reconstruction error without any constraints on the latent space. Such models can capture complex nonlinear mixing but risk learning degenerate solutions where endmembers are neither physically meaningful nor consistent across scenes [12]. At the other extreme, physics-informed architectures incorporate explicit mixing models (e.g., linear or bilinear) as inductive biases, often by structuring the decoder as a linear combination of learnable basis spectra with nonnegative and sum-to-one constraints enforced via softmax or projected gradient layers [13].

The choice of bottleneck dimensionality is a critical architectural parameter. A bottleneck that is too small forces the network to compress spectral variability into a limited number of endmembers, potentially omitting minority materials; a bottleneck that is too large may lead to overfitting and produce redundant or noisy endmembers [14]. Adaptive methods that dynamically select the number of endmembers per scene—or per local region—offer a middle ground, but introduce additional complexity in training and inference. This is especially relevant for large-scale hyperspectral missions such as NASA’s EMIT or ESA’s CHIME, where scenes may contain diverse geological and vegetative compositions.

Spatial context integration is another design dimension. Early joint models processed each pixel independently, but CNNs with encoder-decoder structures—often using U-Net-like skip connections—can propagate spatial information from neighboring pixels to stabilize abundance estimates and smooth spatial transitions [15]. However, spatial aggregation may blur fine-scale boundaries and degrade detection of sub-pixel targets. Attention mechanisms, including self-attention and cross-attention, offer a flexible way to weigh the contribution of

each spatial location and spectral band, but at the cost of quadratic complexity in pixel count. For satellite platforms with limited compute, efficient transformer variants (e.g., shifted windows) or hybrid CNN-transformer designs become necessary.

From a system architecture perspective, the joint model must be embedded in a larger processing pipeline that includes data calibration, dimensionality reduction, atmospheric correction, and post-processing such as spectral library matching. The coupling between these stages and the joint unmixing module can lead to error propagation if not carefully orchestrated. For instance, atmospheric correction algorithms that assume a fixed set of endmembers may be incompatible with a learned representation that adapts to scene-specific variability. Infrastructure designers must therefore consider the entire data lifecycle, from sensor collection to final abundance maps.

4. Robustness and Resilience

Real-world hyperspectral data suffer from multiple sources of degradation: sensor noise, spectral misalignment, atmospheric absorption, and missing bands due to sensor failure or cloud cover. Joint endmember representation models must be robust to these perturbations to ensure reliable abundance maps. One approach is to train the network under a variety of noise models—Gaussian, Poisson, stripe noise—using data augmentation or adversarial training [16]. Attention mechanisms that learn to weigh clean bands more heavily can improve resilience against band-specific artifacts. The weak-signal attention fusion proposed in [18] exemplifies how targeted architectural innovations can amplify signal from rare endmembers even when their spectral signature is partially occluded.

Another dimension of robustness is generalization across different sensors or spatial resolutions. A joint model trained on data from one hyperspectral imager (e.g., AVIRIS-NG) may perform poorly when applied to a different sensor with distinct spectral response functions and signal-to-noise ratios. Domain adaptation techniques, such as adversarial alignment of feature distributions or spectral transfer functions, are necessary to bridge this gap without retraining from scratch [17]. The cost of collecting labeled abundance data for every new sensor is prohibitive, so unsupervised or self-supervised adaptation is highly desirable.

Resilience also involves graceful degradation under adversarial attacks. In defense or security applications, an adversary could inject subtle perturbations to the hyperspectral measurements to cause misclassification of materials (e.g., hiding a military vehicle under a false spectral signature). Joint models that rely heavily on global latent representations may be more susceptible to such attacks than pointwise linear unmixing. Incorporating adversarial training and certified robustness guarantees into the model design—though computationally expensive—is an important research direction.

5. Deployment and Sustainability

Deploying joint unmixing models on orbiting satellites or unmanned aerial vehicles imposes strict constraints on computational resources, memory, and energy consumption. Traditional linear unmixing can be executed in real-time on board with minimal power, while deep joint models often require graphics processing units or dedicated neural accelerators [19]. For large-scale missions, a tiered architecture may be adopted: an onboard lightweight linear unmixing module provides rough abundance estimates for real-time decision-making (e.g., fire detection), while a more sophisticated joint model runs on the ground with full post-

processing resources. The communication bandwidth between satellite and ground is limited, so compressed representations of the latent space could be transmitted for further analysis.

Sustainability of the model itself is a growing concern. Training deep neural networks on large hyperspectral datasets consumes significant energy, especially when hyperparameter optimization and ablation studies are conducted repeatedly. The carbon footprint of each experiment should be reported alongside accuracy metrics, and model compression techniques (pruning, quantization, knowledge distillation) can reduce inference energy without major loss of fidelity [20]. Moreover, the reuse of pretrained models across multiple missions can amortize the training cost, but may introduce biases if the pretraining dataset does not represent the diversity of global land cover.

Data governance also falls under sustainability: hyperspectral datasets are often collected by government agencies or commercial entities with restricted access. Joint models trained on proprietary data may not be reproducible or auditable by third parties, raising concerns about scientific validity and fairness. Open benchmarks such as the Pavia University or Salinas scenes are widely used, but they are limited in size and geographic diversity. Establishing cooperative data-sharing frameworks—similar to the Copernicus program’s free data policy—would accelerate the development of more robust and equitable models.

6. Fairness, Accountability, and Policy Implications

Abundance maps derived from joint endmember and abundance mapping are increasingly used in high-stakes applications: precision agriculture (variable-rate fertilizer application), mineral exploration (drilling permits), urban planning (land-use classification), and environmental enforcement (detecting illegal mining or deforestation). Errors in these maps can have disproportionate impacts on vulnerable communities. For example, if a model systematically underestimates the abundance of a rare indicator mineral in a region inhabited by indigenous groups, it might lead to undervaluation of land rights or missed contamination signals. Conversely, overestimation could result in unwarranted environmental fines.

Fairness in unmixing models can be assessed by examining performance across different geographic regions, land cover types, and spectral conditions. Joint models that rely on neural networks may learn spurious correlations, such as associating high abundance of a certain material with a particular soil type due to regional training data biases [21]. Auditing the latent representations for biases—for instance, using fairness metrics like equalized odds or demographic parity—is challenging because the ground truth for endmembers and abundances is rarely available at scale. Synthetic data generation and simulation-based testing offer partial solutions.

Accountability mechanisms must be built into the governance of systems that deploy these models. When an abundance map leads to an incorrect decision—such as a denied mining permit or a false-positive detection of prohibited crops—there must be a clear chain of responsibility from the model developer to the operator to the regulatory body. Explainability tools, such as attribution maps showing which spectral bands contributed most to a given abundance estimate, can help stakeholders understand and contest outcomes [22]. However, current post-hoc explanations for deep joint models are often unreliable; inherently interpretable architectures (e.g., linear decoders with learnable bases) may be preferable in regulated domains.

Policy considerations extend to international agreements on remote sensing data use. The United Nations Remote Sensing Principles encourage free exchange of data for environmental

monitoring, but commercial hyperspectral constellations are expanding rapidly. Joint unmixing models that are proprietary and trained on private data could create information asymmetries, where only well-resourced actors have access to high-quality abundance maps. Governments and international bodies should consider open-source standards for unmixing algorithms and validation protocols to ensure equitable access and transparency.

7. Conclusion

Joint endmember representation and abundance mapping represents a significant evolution in hyperspectral unmixing, enabling richer integration of spectral-spatial information, nonlinear mixture modeling, and end-to-end optimization. However, the transition from academic proof-of-concept to operational deployment demands a holistic systems perspective that encompasses architectural trade-offs, robustness to real-world corruptions, computational sustainability, and socio-technical governance. This paper has examined these dimensions, highlighting the necessity of aligning model design with infrastructure constraints and fairness imperatives. The weak-signal attention fusion framework [18] exemplifies how targeted architectural innovations can address specific failure modes, but much work remains to standardize evaluation protocols, develop efficient on-board versions, and embed accountability into the deployment pipeline. Future research should prioritize the creation of large-scale, diverse benchmark datasets with pixel-level abundance labels, the development of physics-constrained neural architectures that are both expressive and interpretable, and the formulation of policy guidelines that ensure these powerful analytical tools serve society equitably. Only by addressing systemic challenges can joint endmember and abundance mapping fulfill its promise for precision earth observation.

References

1. Bioucas-Dias, J. M., Plaza, A., Dobigeon, N., Parente, M., Du, Q., Gader, P., & Chanussot, J. (2012). Hyperspectral unmixing overview: Geometrical, statistical, and sparse regression-based approaches. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 5(2), 354-379.
2. Keshava, N., & Mustard, J. F. (2002). Spectral unmixing. *IEEE Signal Processing Magazine*, 19(1), 44-57.
3. Plaza, A., Benediktsson, J. A., Boardman, J. W., Brazile, J., Bruzzone, L., Camps-Valls, G., ... & Trianni, G. (2009). Recent advances in techniques for hyperspectral image processing. *Remote Sensing of Environment*, 113, S110-S122.
4. Dobigeon, N., Tits, L., Somers, B., & Chanussot, J. (2014). A comparison of nonlinear mixing models for hyperspectral image unmixing. *IEEE Geoscience and Remote Sensing Magazine*, 2(3), 28-42.
5. Palsson, B., Sigurdsson, J., Sveinsson, J. R., & Ulfarsson, M. O. (2017). Hyperspectral unmixing using a neural network autoencoder. *IEEE Access*, 6, 25830-25842.
6. Yokoya, N., Chan, J. C., & Segl, K. (2010). Potential of resolution-enhanced hyperspectral data for mineral mapping. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 3(1), 48-58.
7. Winter, M. E. (1999). N-FINDR: An algorithm for fast autonomous spectral end-member determination in hyperspectral data. *Proceedings of SPIE*, 3753, 266-277.

8. Heylen, R., Parente, M., & Gader, P. (2014). A review of nonlinear hyperspectral unmixing methods. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 7(6), 1844-1868.
9. Zhang, Y., Du, B., & Zhang, L. (2018). A spatially adaptive deep learning framework for hyperspectral image unmixing. *IEEE Transactions on Geoscience and Remote Sensing*, 57(6), 3541-3554.
10. Hong, D., Gao, L., Yao, J., Yokoya, N., Chanussot, J., & Zhang, B. (2021). Endmember-guided unmixing network for hyperspectral images. *IEEE Transactions on Geoscience and Remote Sensing*, 60, 1-15.
11. Sobral, L., & Cheung, N. M. (2020). Variational autoencoder for hyperspectral unmixing with uncertainty estimation. *Proceedings of the IEEE International Conference on Image Processing*, 132-136.
12. Rasti, B., Hong, D., Hang, R., Ghamisi, P., Kang, X., Chanussot, J., & Benediktsson, J. A. (2020). Feature extraction for hyperspectral imagery: The evolution from shallow to deep. *IEEE Geoscience and Remote Sensing Magazine*, 8(4), 48-78.
13. Zare, A., & Ho, K. C. (2014). Endmember variability in hyperspectral analysis. *IEEE Signal Processing Magazine*, 31(1), 95-104.
14. Jiang, J., Ma, J., Wang, Z., Chen, C., & Wang, L. (2020). Hyperspectral unmixing via deep autoencoder with nonnegative and sum-to-one constraints. *IEEE Geoscience and Remote Sensing Letters*, 17(9), 1578-1582.
15. Xu, J., Li, S., & Sun, X. (2021). A spatial-spectral U-Net for hyperspectral unmixing. *IEEE Transactions on Image Processing*, 30, 7422-7435.
16. Camps-Valls, G., Tuia, D., Bruzzone, L., & Benediktsson, J. A. (2014). Advances in hyperspectral image classification: Earth monitoring with statistical learning methods. *IEEE Signal Processing Magazine*, 31(1), 45-54.
17. Tuia, D., Camps-Valls, G., & Benediktsson, J. A. (2011). Domain adaptation for the classification of remote sensing data: An overview. *IEEE Geoscience and Remote Sensing Magazine*, 2(2), 18-39.
18. Long, Z., Zia, A., Fu, G., Rolland, V., & Zhou, J. (2026). WS-Net: Weak-Signal Representation Learning and Gated Abundance Reconstruction for Hyperspectral Unmixing via State-Space and Weak Signal Attention Fusion. *arXiv preprint arXiv:2603.09037*.
19. Verhoeven, G., & Doneus, M. (2011). Balancing on the borderline: A low-cost approach to visualize the thermal properties of the ground for archaeological prospection. *Archaeological Prospection*, 18(4), 253-262.
20. Jagannathan, S., & Bhatt, D. (2021). Energy-efficient deep learning for remote sensing: A survey. *IEEE Transactions on Green Communications and Networking*, 5(3), 1101-1115.
21. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), 1-35.
22. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?" Explaining the predictions of any classifier. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1135-1144.