

Knowledge-Guided Masked Autoencoder with Linear Spectral Mixing and Spectral-Angle-Aware Reconstruction

Abdul Matin¹, Rupasree Dey¹, Tanjim Bin Faruk¹, Shrideep Pallickara¹, Sangmi Lee Pallickara¹

¹Department of Computer Science, Colorado State University
Fort Collins, CO, USA

amatin@colostate.edu, rupasree.dey@colostate.edu, tanjim@colostate.edu, shrideep@colostate.edu, sangmi@colostate.edu

Abstract

Integrating domain knowledge into deep learning has emerged as a promising direction for improving model interpretability, generalization, and data efficiency. In this work, we present a novel knowledge-guided ViT based Masked Autoencoder that embeds scientific domain knowledge within the self-supervised reconstruction process. Instead of relying solely on data-driven optimization, our proposed approach incorporates the Linear Spectral Mixing Model (LSMM) as a physical constraint and physically-based Spectral Angle Mapper (SAM) ensuring that learned representations adhere to known structural relationships between observed signals and their latent components. The framework jointly optimizes LSMM and SAM loss with a conventional Huber loss objective, promoting both numerical accuracy and geometric consistency in the feature space. This knowledge-guided design enhances reconstruction fidelity, stabilizes training under limited supervision, and yields interpretable latent representations grounded in physical principles. The experimental findings indicate that the proposed model substantially enhances reconstruction quality and improves downstream task performance, highlighting the promise of embedding physics-informed inductive biases within transformer-based self-supervised learning.

Introduction

Recent advances in self-supervised learning and large-scale transformer architectures have demonstrated impressive representation learning capabilities across vision and multimodal tasks. In particular, Vision Transformer-based Masked Autoencoders (ViT-MAEs) have shown that reconstructing masked input patches provides an effective pretext task for learning transferable features at scale (He et al. 2022; Dosovitskiy et al. 2021). Such architectures have further evolved into large-scale *foundation models* that can adapt across diverse visual and geospatial applications by leveraging vast unlabeled data (Bommasani et al. 2021). However, purely data-driven reconstruction objectives often overlook domain-specific priors and the physical mechanisms that govern data generation, limiting interpretability and generalization (Willard et al. 2020; Raissi, Perdikaris, and Karniadakis 2019).

In the context of Earth observation, hyperspectral satellite imagery offers a unique and information-rich modality. By capturing hundreds of contiguous spectral bands, these data provide a fine-grained description of surface materials, enabling precise detection of land-cover types, vegetation health, mineral composition, and soil properties (Zhang et al. 2023). Despite their potential, hyperspectral datasets present distinct challenges for deep learning models: high spectral dimensionality, spatial redundancy, and the physical phenomenon of mixed pixels where each observation represents a linear or nonlinear combination of multiple materials. Consequently, the success of general-purpose foundation models cannot be directly transferred to hyperspectral data without incorporating the underlying physical knowledge that explains such mixtures.

In recent years, an emerging paradigm known as *knowledge-guided machine learning* (KGML) has begun to blur the boundary between purely data-driven methods and science-driven modeling. Rather than treating machine learning as a black box, KGML frameworks strategically integrate domain knowledge (such as physical laws, mixing models, or causally grounded relationships) with representation learning in neural architectures (Karpadne, Jia, and Kumar 2024). For example, Liu et al. (2024) demonstrate how embedding process-based and remote-sensing knowledge into neural frameworks significantly improves soil-carbon estimation across agroecosystems, highlighting the potential of hybrid scientific-data approaches. Drawing on these insights, our work situates itself at the intersection of foundation-model design and domain-informed representation learning.

Motivated by these observations, we propose the **Knowledge-Augmented Reconstruction with Masked Autoencoding (KARMA)**, a physics-informed ViT-MAE framework that fuses a physically interpretable model with transformer-based self-supervision. KARMA integrates the Linear Spectral Mixing Model (LSMM) (Bioucas-Dias et al. 2012) as an inductive bias inside the decoder: the decoder predicts an abundance vector per patch, and the model reconstructs the observed signal as a linear combination of surface material signatures. This physics-guided structure ensures that learned latent representations remain consistent with real-world spectral composition laws, thereby promoting explainability and stability.

To further promote spectral fidelity and geometric alignment in the learned features, we jointly optimize a Spectral-Angle-aware loss along with Huber Loss and an explicit LSMM physics consistency loss. This hybrid objective encourages numerical accuracy, shape-preserving reconstructions, and explicit physical consistency which results in producing interpretable latent components that are useful for downstream tasks such as regression and classification. By embedding physical priors into a ViT-MAE backbone, KARMA bridges the gap between generic self-supervised learning and domain-grounded scientific understanding—offering a step toward foundation models that learn not only from data but also from the governing physical principles.

Our contributions include the following:

- We introduce **KARMA**, a ViT-MAE variant that integrates the Linear Spectral Mixing Model (LSMM) as a physics-guided branch and enforces abundance constraints in an end-to-end manner.
- We design a hybrid objective combining Huber loss, Spectral Angle Mapper (SAM), and a physics-consistency term that jointly improves spectral fidelity and physical realism.
- We demonstrate the effectiveness of KARMA on the EnMAP hyperspectral benchmark, showing that physics-guided learning improves reconstruction fidelity, interpretability, and downstream transfer across both regression and classification tasks.

Methodology

This section presents the proposed KARMA framework. We begin with an overview of the ViT-MAE backbone architecture, then elaborate on our physics-informed learning with LSMM, spectral constraint, SAM and describe the hybrid optimization objective. The complete proposed architecture is depicted in Figure 1.

ViT-MAE backbone

The base architecture of KARMA builds upon the asymmetric Vision Transformer-based Masked Autoencoder (ViT-MAE) (He et al. 2022; Dosovitskiy et al. 2021). The input hyperspectral cube is divided into non-overlapping patches, where each of the patches are flattened and linearly projected into an embedding vector. Following standard practice, a high masking ratio (typically 75%) is applied to randomly mask a large portion of these patches, creating a reconstruction-based self-supervised task that encourages the model to learn robust contextual dependencies.

The encoder processes only the visible (unmasked) patches, producing compact latent representations that summarize spatial-spectral context. The lightweight decoder then reconstructs the masked patches using both the encoded latent tokens and learnable mask tokens. Unlike standard MAE architectures that predict pixel values directly, our decoder incorporates a physics-guided branch that explicitly models the physical composition of each spectral observation, introducing domain knowledge directly into the learning process.

Integration of Linear Spectral Mixing Model (LSMM)

To bridge physical modeling with transformer-based learning, we embed the LSMM into the decoder as a parallel reconstruction branch. LSMM is a physics-based model in remote sensing and spectral analysis which is used to describe how the spectral signature recorded by a sensor over a pixel is formed by the combination of different surface materials (endmembers) (Wei and Wang 2020). LSMM assumes that an observed signal can be represented as a linear combination of those latent components or *endmembers*:

$$\mathbf{r} = \mathbf{A}\mathbf{x} + \mathbf{e}, \quad (1)$$

where $\mathbf{r} \in \mathbb{R}^C$ is the observed spectral vector of EnMap hyperspectral images with $C = 218$ bands, $\mathbf{A} \in \mathbb{R}^{C \times M}$ is the endmember matrix containing the spectral signatures of M constituent materials, $\mathbf{x} \in \mathbb{R}^M$ is the abundance vector, and \mathbf{e} denotes reconstruction noise (Bioucas-Dias et al. 2012).

Physical validity requires that abundances satisfy two fundamental constraints, *sum-to-one* and *non-negativity*:

$$\mathbf{x} \geq 0, \quad \mathbf{1}^\top \mathbf{x} = 1, \quad (2)$$

These constraints ensure that the material mixtures remain physically meaningful. In our KARMA architecture, the decoder incorporates a compact *abundance head* $f_\theta(\cdot)$, implemented as a lightweight multilayer perceptron (MLP). This module transforms each decoder patch token \mathbf{z} into predicted abundance values. By applying softmax activation, we obtain

$$\hat{\mathbf{x}} = \text{softmax}(f_\theta(\mathbf{z})), \quad (3)$$

which inherently enforces both the non-negativity and simplex constraints without the need for additional regularization.

The LSMM-based reconstruction is then computed as:

$$\hat{\mathbf{r}}_{\text{phys}} = \mathbf{A}\hat{\mathbf{x}}. \quad (4)$$

The endmember matrix \mathbf{A} can be initialized from prior knowledge (e.g., spectral libraries or k -Means prototypes) or learned end-to-end with a small learning rate for adaptive refinement (Ozkan, Kaya, and Akar 2019).

Spectral Angle Mapper (SAM) for geometric alignment

The Spectral Angle Mapper (SAM) offers a complementary perspective to traditional pixel-wise metrics by measuring the angular distance between spectral vectors, independent of their magnitude (Kruse et al. 1993; Yuhas, Goetz, and Boardman 1992). Unlike MSE or Huber Loss, which is sensitive to intensity variations, SAM focuses on preserving the spectral *shape*, a critical property for identifying materials in hyperspectral data. We define the SAM loss across N samples as:

$$\mathcal{L}_{\text{SAM}} = \frac{1}{N} \sum_{i=1}^N \arccos \left(\frac{\langle \hat{\mathbf{r}}_i, \mathbf{r}_i \rangle}{\|\hat{\mathbf{r}}_i\|_2 \|\mathbf{r}_i\|_2 + \epsilon} \right), \quad (5)$$

where $\langle \cdot, \cdot \rangle$ denotes the dot product, and ϵ is a small constant that prevents numerical instability from division by zero.

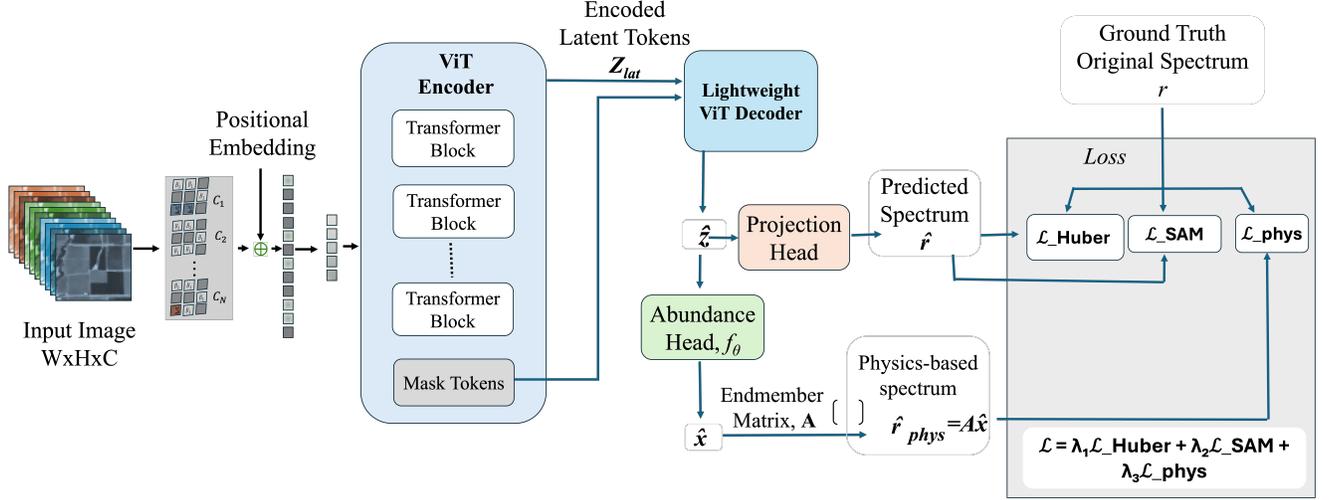


Figure 1: Architecture of KARMA: A knowledge-guided ViT-MAE that integrates the Linear Spectral Mixing Model (LSMM) within the decoder. The model jointly optimizes Huber loss, SAM, and physics-consistency losses to enhance spectral fidelity and interpretability.

Minimizing \mathcal{L}_{SAM} guides the reconstructed spectra $\hat{\mathbf{r}}$ to align directionally with the ground truth \mathbf{r} in spectral space, even when absolute magnitudes differ. This geometric alignment helps the model maintain consistent representations and enhances its ability to discriminate between spectrally similar materials.

Hybrid objective function

KARMA optimizes three complementary objectives that jointly enforce numerical accuracy, spectral shape fidelity, and physical consistency. The unified loss function combines:

$$\mathcal{L} = \lambda_1 \mathcal{L}_{\text{Huber}} + \lambda_2 \mathcal{L}_{\text{SAM}} + \lambda_3 \mathcal{L}_{\text{phys}}, \quad (6)$$

where λ_i are hyperparameters that denote the weights to control the contribution of each term.

(1) Huber Loss (Robust Regression). We employ the *Huber loss* to improve robustness against spectral outliers and noise. The Huber loss behaves quadratically for small errors and linearly for large deviations, thereby balancing sensitivity and stability during reconstruction:

$$\mathcal{L}_{\text{Huber}} = \frac{1}{N} \sum_{i=1}^N \begin{cases} \frac{1}{2} \|\hat{\mathbf{r}}_i - \mathbf{r}_i\|_2^2, & \text{if } \|\hat{\mathbf{r}}_i - \mathbf{r}_i\|_2 \leq \delta, \\ \delta \|\hat{\mathbf{r}}_i - \mathbf{r}_i\|_2 - \frac{1}{2} \delta^2, & \text{otherwise,} \end{cases} \quad (7)$$

(2) Spectral Angle Mapper (SAM). Enforces angular similarity and shape preservation, as defined above.

(3) Physics Consistency Loss. Ensures reconstruction fidelity constrained by LSMM:

$$\mathcal{L}_{\text{phys}} = \frac{1}{N} \sum_{i=1}^N \|\mathbf{r}_i - \mathbf{A} \hat{\mathbf{x}}_i\|_2^2. \quad (8)$$

Together, these losses create a balance between data-driven and knowledge-driven supervision, allowing the model to learn both expressive and interpretable representations.

Architectural details

Our architecture begins with patch embedding that projects the $C = 218$ input channels (from the EnMap hyperspectral satellite images) into an embedding dimension D (typically $D = 512$). The multi-head self-attention mechanism employs H heads with D divisible by H (commonly $H = 8$). The abundance head comprises a compact MLP with architecture $D \rightarrow D/2 \rightarrow M$, producing abundance predictions for each patch token. We set the endmember matrix \mathbf{A} dimensions to $218 \times M$, where M typically based on dataset characteristics and ablation studies like as shown in Figure 2

This physics-guided training regulates KARMA to ensure that reconstructions are both numerically accurate and physically consistent. This is an essential step toward interpretable foundation models for hyperspectral data.

Experimental Setup

Dataset and preprocessing

We conduct our experiments using hyperspectral imagery from the Environmental Mapping and Analysis Program (EnMAP) (Storch et al. 2023), a spaceborne sensor that delivers observations across 224 spectral bands spanning the 420–2450 nm frequency range at 30 m spatial resolution. Our primary study region covers the state of California, which offers rich diversity in land-cover types ranging from dense forests and urban areas to heterogeneous agricultural landscapes. To test the generalization of our model across different environmental conditions, we expand our downstream analysis to include representative areas in Colorado and Kansas.

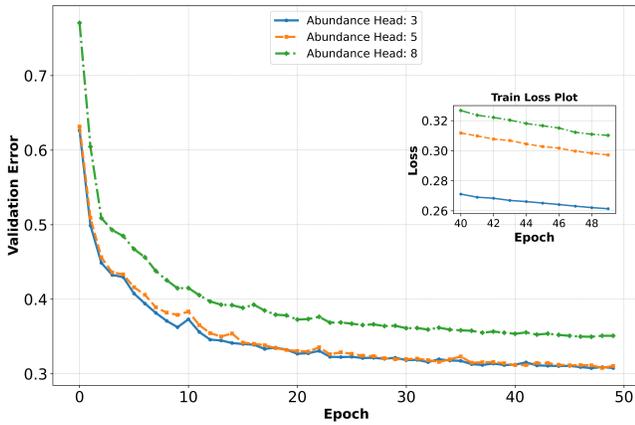


Figure 2: Ablation study on the number of abundance head components. The results show that using three abundance heads achieves faster convergence and stable reconstruction performance while maintaining lower computational cost.

Following radiometric calibration and atmospheric correction, we removed six bands affected by noise and retained 218 spectral bands for our analysis. We spatially resample the imagery into non-overlapping 224×224 pixel tiles, matching the input requirements of our ViT-MAE-based architecture. The dataset is organized into three subsets: 5,000 tiles for pretraining, 500 for validation, 2,000 for the primary test set (*Test Dataset 1*). We conducted certain ablation studies on smaller subsets of these data due to computational constraints.

Downstream Model

As illustrated in Figure 3, we leverage the pre-trained **KARMA encoder** to extract rich spectral-spatial representations from the 218-band EnMAP hyperspectral data. For downstream tasks, lightweight task-specific heads are attached to the frozen encoder to ensure efficient fine-tuning while preserving the learned representations. Here, we employ a compact classification head composed of a few convolutional and fully connected layers, optimized using the cross-entropy loss for two target applications.

Training details

We adopt the following default configuration: 16×16 patch size, embedding dimension $D = 512$, $H = 8$ attention heads, 75% masking ratio, and batch size of 32. Training uses the AdamW optimizer with a base learning rate of 1×10^{-4} and cosine decay schedule. Pretraining runs for 300-500 epochs depending on computational resources. The hybrid loss weight parameters, λ_1 , λ_2 , and λ_3 are also chosen empirically. The endmember matrix \mathbf{A} is initialized randomly in this work.

Evaluation protocol

Our evaluation follows a two-phase structure designed to assess both reconstruction quality and transferability of learned representations.

Phase I: Reconstruction Quality. We evaluate reconstruction performance using PSNR and SSIM. While PSNR assesses pixel-level fidelity, indicating how closely the reconstructed spectra match the reference, SSIM provides a perceptual measure of structural integrity, reflecting the model’s ability to preserve spatial and textural patterns.

Phase II: Downstream Transfer Learning. We utilize the pretrained KARMA encoder across two representative downstream tasks to assess the quality and generalization of its learned representations. The first task, crop type identification, is formulated as a multi-class classification problem spanning diverse agricultural categories. The second task, national land cover classification, captures broader spatial and ecological variability. To evaluate performance, we report both Top-1 Accuracy and Mean Intersection over Union (mIoU), which together measure categorical precision and spatial consistency in the model’s predictions.

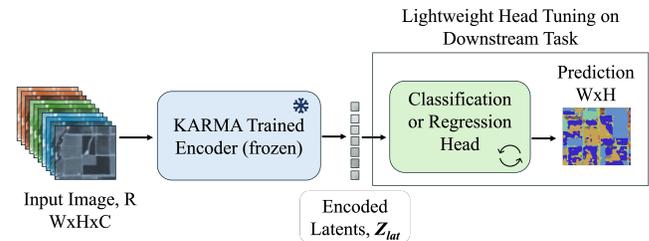


Figure 3: Downstream model consisting of the frozen KARMA encoder (pre-trained with KG) and a lightweight CNN head

Results and Analysis

Reconstruction Fidelity Analysis

Table 1 demonstrates that incorporating physical knowledge through the Linear Spectral Mixing Model (LSMM) and Spectral Angle Mapper (SAM) loss improves reconstruction quality over the purely data-driven baseline. KARMA achieves an average PSNR of 27.38 dB compared to ViT-MAE’s 24.61 dB, representing a 2.77 dB improvement. This corresponds to an 11.3% relative gain in reconstruction accuracy.

The improvement is more pronounced when measured by the Structural Similarity Index Measure (SSIM), which captures structural pattern preservation beyond pixel-wise differences. KARMA achieves an average SSIM of 0.68 compared to ViTMAE’s 0.55, a relative improvement of 23.6%. The higher SSIM scores indicate that KARMA produces reconstructions that preserve physically meaningful spectral-spatial structures, validating the incorporation of domain-specific physical constraints into the masked autoencoder framework.

Downstream Task Performance and Generalization

While the classification of certain classes in CDL and NLCD is inherently challenging due to noisy labels and dataset preparation issues, the improvements in self-supervised reconstruction translate into significant gains across down-

ID	Model	Reconstruction Loss	Physical Constraints	Avg. PSNR (\uparrow)	Per-Channel PSNR (Max) (\uparrow)	Avg. SSIM (\uparrow)	Per-Channel SSIM (Max) (\uparrow)
1	VitMAE	HUBER	-	24.61	32.45	0.55	0.82
2	KARMA (Ours)	HUBER	SAM + LSMM	27.38	35.10	0.68	0.88
Improvement (% over ViTMAE)				11.26	8.17	23.64	7.32

Table 1: Reconstruction accuracy comparison between KARMA and ViTMAE on hyperspectral data. The physics-guided learning in KARMA consistently yields superior reconstruction fidelity.

Class Name	ViTMAE + Head		KARMA + Head	
	Top-1 Acc (%) (\uparrow)	mIoU (%) (\uparrow)	Top-1 Acc (%) (\uparrow)	mIoU (%) (\uparrow)
Shrubland	40.31	34.74	71.18	51.35
Grassland/Pasture	22.74	16.85	44.22	31.47
Evergreen Forest	81.72	53.04	85.03	56.30
Average	48.26	34.88	66.81	46.37

Table 2: Crop Type Identification on CDL data (top three classes in California). KARMA with a lightweight CNN significantly outperforms the baseline.

stream tasks. This highlights the interpretability and transferability of KARMA’s learned features.

Crop Type Identification (CDL Data) As reported in Table 2, KARMA paired with a lightweight classification head substantially outperforms ViTMAE. On average, Top-1 Accuracy rises from 48.25% to 63.12% (a 30.8% relative improvement), and mean IoU increases from 35.67% to 50.88% (a 42.6% relative gain). The largest improvements are observed in spectrally complex and mixed classes, demonstrating that LSMM-guided pre-training produces discriminative and physically grounded latent features despite dataset noise.

Land Cover Classification (NLCD Data) Table 3 evaluates generalization by testing the KARMA encoder, trained on California data, on geographically distinct regions (Colorado and Kansas). KARMA consistently achieves higher accuracy across all classes, with the most notable improvement in "Cultivated Crops," rising from 56.70% to 91.59% (a 61.5% relative increase). Overall, average Top-1 Accuracy improves from 61.24% to 79.83% (+30.3% relative) and mean IoU from 44.18% to 63.47% (+43.6% relative). These results confirm that KARMA captures physics-based inductive biases, yielding features that are robust and transferable across noisy datasets and diverse terrains.

Related Work

Self-supervised vision and masked modeling. MAEs and related masked modeling approaches have rapidly advanced representation learning for vision, showing strong transfer to downstream tasks (He et al. 2022; Dosovitskiy

Class Name	Top-1 Accuracy (%) (\uparrow)	
	ViTMAE + Head	KARMA + Head
Shrub/Scrub	63.18	70.50
Herbaceous	35.43	46.59
Cultivated Crops	56.70	91.59
Average	51.77	69.56

Table 3: Land Cover Classification on NLCD data (top three classes across CA, CO, and KS). The KARMA encoder, trained solely on California data, generalizes effectively to other regions (Colorado and Kansas), demonstrating improved downstream performance.

et al. 2021). Extensions to multi-spectral and temporal remote sensing (e.g., SatMAE) demonstrate benefits from spectral-temporal pretext tasks (Cong et al. 2022). Building on this direction *HyperKD* (Matin et al. 2025), extended the ViT-MAE paradigm to cross-spectral settings through inverse domain knowledge distillation, effectively bridging multispectral and hyperspectral modalities. In contrast, the present study departs from distillation-based strategies and instead introduces a physics-guided mechanism that embeds the Linear Spectral Mixing Model (LSMM) directly within the ViT-MAE architecture to enhance reconstruction fidelity and interpretability.

Physics-informed and knowledge-guided learning. Integrating scientific priors into ML models has been formalized under physics-informed learning and knowledge-guided ML; these approaches provide improved physical realism and extrapolation (Raissi, Perdikaris, and Karniadakis 2019; Willard et al. 2020; Karpatne, Jia, and Kumar 2024; Liu et al. 2024). KGML methods have been used to incorporate support for physical constraints grounded in soil hydrology, including the van Genuchten water retention equations and models (i.e., Richards’ Equation) of hydraulic conductivity (Khandelwal, Pallickara, and Pallickara 2024; Khandelwal et al. 2025); vegetation indices (Bruhwiler et al. 2020); evapotranspiration (Armstrong et al. 2022); preserving graph properties such as betweenness centrality (Matin et al. 2022); masking cloud occlusions in satellite imagery (Khandelwal et al. 2022), accounting for human perceptual limits during visualization (Mitra et al. 2021); masked autoencoders (Faruk et al. 2025; Matin et al. 2025) and ac-

counting for correlations between soil spectroscopic properties (Bachinin et al. 2025; Dey et al. 2025). Our KARMA framework follows this paradigm by embedding LSMM within a self-supervised transformer framework. Finally, knowledge distillation approaches have been explored in the context of physical phenomena (Matin et al. 2023); our methodology is complementary to these approaches.

Spectral mixing and unmixing. LSMM and hyperspectral unmixing have a rich literature; classical formulations and modern deep unmixing approaches show how material signatures can be recovered and used for interpretability (Bioucas-Dias et al. 2012; Ozkan, Kaya, and Akar 2019). KARMA uses LSMM as an explicit constraint rather than a standalone unmixing network.

Conclusion and Future Work

We introduced KARMA, a knowledge-guided masked autoencoder that bridges data-driven learning with physical principles for hyperspectral image analysis. By integrating the Linear Spectral Mixing Model into the decoder and using a hybrid loss that combines Huber, SAM, and physics-based constraints, our approach yields more accurate and physically interpretable reconstructions. Our experiments with EnMAP hyperspectral imagery demonstrates that this physics-guided self-supervision strategy not only improves reconstruction fidelity but also yields representations that transfer more effectively to downstream tasks. Our work adds to the growing body of evidence showing that incorporating domain theoretic knowledge into modern self-supervised architectures leads to models that achieve strong performance while remaining scientifically meaningful. This is particularly crucial for remote sensing applications where physical consistency matters as much as predictive performance.

As we move forward, we plan to broaden our study evaluating KARMA against several established baselines. Complementary ablation studies will systematically analyze the contribution of each component—such as the SAM loss, LSMM integration, the role of fixed versus learnable endmember matrices A , and the effect of varying the number of endmembers M . Beyond these evaluations, we plan to explore scalability and cross-domain generalization by combining hyperspectral data with complementary sensing modalities. Together, these directions will strengthen the bridge between physical understanding and machine learning, ultimately leading to more robust, reliable, and interpretable Earth observation models.

Acknowledgments

This research was supported by the National Science Foundation (1931363,2312319), the National Institute of Food Agriculture (COL014021223, 2025-77039-45531), and an NSF/NIFA Artificial Intelligence Institutes AI-LEAF Award [2023-03616].

References

Armstrong, S.; Khandelwal, P.; Padalia, D.; Senay, G.; Schulte, D.; Andales, A.; Breidt, F. J.; Pallickara, S.; and Pal-

lickara, S. L. 2022. Attention-based convolutional capsules for evapotranspiration estimation at scale. *Environmental Modelling & Software*, 152: 105366.

Bachinin, A.; Dey, R.; Khandelwal, P.; Leuthold, S.; Cotrufo, M. F.; Pallickara, S.; and Pallickara, S. L. 2025. Science-Informed Multitask Transformer for Soil Property Prediction from FTIR Spectroscopy. In *2025 IEEE International Conference on eScience (eScience)*, 48–57. IEEE.

Bioucas-Dias, J. M.; Plaza, A.; Dobigeon, N.; Parente, M.; Du, Q.; Gader, P.; and 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.

Bommasani, R.; Hudson, D. A.; Adeli, E.; Altman, R.; Arora, S.; von Arx, S.; Bernstein, M. S.; Bohg, J.; Bosselut, A.; Brunskill, E.; Brynjolfsson, E.; et al. 2021. On the Opportunities and Risks of Foundation Models. *arXiv preprint arXiv:2108.07258*.

Bruhwieler, K.; Khandelwal, P.; Rammer, D.; Armstrong, S.; Pallickara, S. L.; and Pallickara, S. 2020. Lightweight, embeddings based storage and model construction over satellite data collections. In *2020 IEEE International Conference on Big Data (Big Data)*, 246–255. IEEE.

Cong, Y.; Khanna, S.; Meng, C.; Liu, P.; Rozi, E.; He, Y.; Burke, M.; Lobell, D. B.; and Ermon, S. 2022. SatMAE: Pre-training Transformers for Temporal and Multi-Spectral Satellite Imagery. In *Advances in Neural Information Processing Systems (NeurIPS) Workshops*.

Dey, R.; Matin, A.; Lewark, E.; Faruk, T. B.; Bachinin, A.; Leuthold, S.; Cotrufo, M. F.; Pallickara, S.; and Pallickara, S. L. 2025. DeepSalt: Bridging Laboratory and Satellite Spectra through Domain Adaptation and Knowledge Distillation for Large-Scale Soil Salinity Estimation. In *IEEE International Conference on Big Data 2025*.

Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Houlsby, N. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations (ICLR)*.

Faruk, T. B.; Matin, A.; Pallickara, S.; and Pallickara, S. L. 2025. Accounting for Spatial Variability with the Histogram of Oriented Gradients Based Masking Improves Performance of Masked Autoencoder over Hyperspectral Satellite Imagery (Student Abstract). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 29365–29367.

He, K.; Chen, X.; Xie, S.; Li, Y.; Dollár, P.; and Girshick, R. 2022. Masked Autoencoders Are Scalable Vision Learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 16000–16012.

Karpatne, A.; Jia, X.; and Kumar, V. 2024. Knowledge-guided Machine Learning: Current Trends and Future Prospects. *arXiv preprint*, arXiv:2403.15989.

Khandelwal, P.; Armstrong, S.; Matin, A.; Pallickara, S.; and Pallickara, S. L. 2022. Cloudnet: A deep learning approach

- for mitigating occlusions in landsat-8 imagery using data coalescence. In *2022 IEEE 18th International Conference on e-Science (e-Science)*, 117–127. IEEE.
- Khandelwal, P.; Niemann, J. D.; Mulla, D. J.; Pallickara, S.; and Pallickara, S. L. 2025. Subterra: Estimating soil moisture at root zone depths using science-guided learning. In *2025 IEEE Conference on Artificial Intelligence (CAI)*, 328–335. IEEE.
- Khandelwal, P.; Pallickara, S. L.; and Pallickara, S. 2024. Deepsoil: A science-guided framework for generating high precision soil moisture maps by reconciling measurement profiles across in-situ and remote sensing data. In *Proceedings of the 32nd ACM International Conference on Advances in Geographic Information Systems*, 233–246.
- Kruse, F. A.; Lefkoff, A. B.; Boardman, J. W.; Heidebrecht, K. B.; Shapiro, A. T.; Barloon, P. J.; and Goetz, A. F. H. 1993. The Spectral Image Processing System (SIPS)—Interactive visualization and analysis of imaging spectrometer data. *Remote Sensing of Environment*, 44(2–3): 145–163.
- Liu, L.; Zhou, W.; Guan, K.; Peng, B.; Xu, S.; Tang, J.; Zhu, Q.; Till, J.; Jia, X.; Jiang, C.; Wang, S.; Qin, Z.; Kong, H.; Grant, R.; Mezbahuddin, S.; Kumar, V.; and Jin, Z. 2024. Knowledge-guided machine learning can improve carbon cycle quantification in agroecosystems. *Nature Communications*, 15(1): 357.
- Matin, A.; Armstrong, S.; Mitra, S.; Pallickara, S.; and Pallickara, S. L. 2022. Rapid betweenness centrality estimates for transportation networks using capsule networks. In *2022 Fourth International Conference on Transdisciplinary AI (TransAI)*, 89–96. IEEE.
- Matin, A.; Faruk, T. B.; Pallickara, S.; and Pallickara, S. L. 2025. HyperKD: Distilling Cross-Spectral Knowledge in Masked Autoencoders via Inverse Domain Shift with Spatial-Aware Masking and Specialized Loss. In *Proceedings of the IEEE International Conference on Data Science and Advanced Analytics (DSAA)*. Birmingham, United Kingdom.
- Matin, A.; Khandelwal, P.; Pallickara, S.; and Pallickara, S. L. 2023. Discern: Leveraging knowledge distillation to generate high resolution soil moisture estimation from coarse satellite data. In *2023 IEEE International Conference on Big Data (BigData)*, 1222–1229. IEEE.
- Mitra, S.; Rammer, D.; Pallickara, S.; and Pallickara, S. L. 2021. Glance: A generative approach to interactive visualization of voluminous satellite imagery. In *2021 IEEE International Conference on Big Data (Big Data)*, 359–367. IEEE.
- Ozkan, S.; Kaya, B.; and Akar, G. B. 2019. EndNet: Sparse AutoEncoder Network for Endmember Extraction and Hyperspectral Unmixing. *IEEE Transactions on Geoscience and Remote Sensing*, 57(1): 482–496.
- Raissi, M.; Perdikaris, P.; and Karniadakis, G. E. 2019. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378: 686–707.
- Storch, T.; Honold, H.-P.; Chabrilat, S.; Habermeyer, M.; Tucker, P.; Brell, M.; Ohndorf, A.; Wirth, K.; Betz, M.; Kuchler, M.; et al. 2023. The EnMAP imaging spectroscopy mission towards operations. *Remote Sensing of Environment*, 294: 113632.
- Wei, J.; and Wang, X. 2020. An overview on linear unmixing of hyperspectral data. *Mathematical Problems in Engineering*, 2020(1): 3735403.
- Willard, J.; Jia, X.; Xu, S.; Steinbach, M.; and Kumar, V. 2020. Integrating scientific knowledge with machine learning for engineering and environmental systems. *arXiv preprint arXiv:2003.04919*.
- Yuhas, R. H.; Goetz, A. F.; and Boardman, J. W. 1992. Discrimination among semi-arid landscape endmembers using the Spectral Angle Mapper (SAM) algorithm. In *Summaries of the Third Annual JPL Airborne Geoscience Workshop*, 147–149.
- Zhang, Y.; Gao, L.; Zhang, B.; Plaza, A.; and Chanussot, J. 2023. Hyperspectral Remote Sensing for Earth Observation: Recent Advances, Challenges, and Future Trends. *IEEE Geoscience and Remote Sensing Magazine*, 11(2): 8–38.