

# WorldModel-Unmix: Generative Earth Observation World Models for Spectral Evolution Prediction and Hyperspectral Unmixing Analysis

Jerome R. Perkins

Department of Computer Science, University of Central Florida, Orlando, FL, USA.  
jrperkins@ucf.edu

Jose L. Jarvinen

Department of Computer Science, University of Houston, Houston, TX, USA.  
jose.jarvinen@uh.edu

Keith L. Sanders

School of Electrical Engineering and Computer Science, Oregon State University, Corvallis, OR, USA.  
ksanders@oregonstate.edu

Mikko Hawkins

Department of Computer Science and Engineering, University of Nevada, Reno, Reno, NV, USA.  
hellomikko@unr.edu

## Abstract

Earth observation missions have become indispensable for monitoring environmental change, resource management, and global security. Hyperspectral imaging, in particular, provides rich spectral signatures that enable the identification of materials and their abundances across the planet surface. However, the inherent complexity of spectral mixing, temporal dynamics, and the limited availability of ground-truth data pose significant challenges to traditional unmixing methods. This paper introduces WorldModel-Unmix, a generative world model framework designed for earth observation that integrates large-scale spectral evolution prediction with hyperspectral unmixing analysis. The proposed system leverages a learned latent representation of the Earth's surface processes to forecast spectral changes over time while simultaneously decomposing mixed pixels into constituent materials. We discuss the architectural trade-offs between generative capacity and physical consistency, the infrastructure requirements for global deployment, and the socio-technical implications of such a model for environmental governance and fairness. By embedding unmixing within a world model, the framework achieves robust performance under distribution shift and data scarcity, while also enabling counterfactual reasoning about future land cover scenarios. The paper examines deployment considerations, sustainability of large-scale training, and policy challenges related to transparency and accountability. Through cross-domain comparisons with climate modeling and autonomous driving, we highlight the unique demands of earth observation world models. Finally, we outline open research directions including uncertainty quantification, federated learning for global coverage, and alignment with international remote sensing standards. WorldModel-Unmix represents a step toward unified, generative earth intelligence that reconciles prediction and analysis in a single scalable architecture.

## Keywords

Earth observation, hyperspectral unmixing, world models, spectral evolution, generative AI, remote sensing, large-scale systems, socio-technical infrastructure, sustainability, fairness.

## 1. Introduction

The accelerating pace of environmental change demands continuous, high-resolution monitoring of the Earth’s surface. Hyperspectral remote sensing offers unparalleled spectral detail, capturing hundreds of narrow bands that allow the identification of minerals, vegetation types, atmospheric constituents, and anthropogenic materials. Yet the promise of hyperspectral imagery is tempered by the mixed nature of every pixel: most spatial resolutions are too coarse to isolate pure materials, so each measured spectrum is a composite of multiple substances. Hyperspectral unmixing—the process of estimating the fractional abundances of endmembers within a pixel—has been a central research challenge for decades. Traditional approaches range from linear mixture models with geometric constraints to nonlinear formulations that account for multiple scattering. However, these methods often assume static scenes and rely on extensive prior knowledge or spectral libraries that may not generalize across time and geography.

Meanwhile, the rise of generative world models in artificial intelligence has opened new possibilities for understanding and predicting complex dynamical systems. World models learn compressed representations of an environment’s state and transition dynamics, enabling agents to simulate future outcomes and reason about interventions. Applied to Earth observation, a world model could learn the temporal evolution of land surface spectra from historical satellite data, capturing natural cycles, seasonal trends, and abrupt events. By integrating unmixing directly into the generative framework, we can move beyond static decomposition to a dynamic, predictive analysis that accounts for how material abundances change over time. This is the core motivation behind WorldModel-Unmix: a generative architecture that simultaneously performs spectral evolution prediction and hyperspectral unmixing, all within a learned latent space that embodies physical consistency and robust generalization.

In this paper, we present the system-level design of WorldModel-Unmix, focusing on architectural choices, trade-offs, and the broader context of deployment, sustainability, and governance. We argue that such a unified model offers distinct advantages over separate pipelines for prediction and unmixing, including reduced error propagation, better handling of missing data, and the ability to generate counterfactual scenarios for policy planning. However, these benefits come with substantial computational costs and challenges in training data representativeness, model interpretability, and fairness across different regions. The paper is organized as follows. Section 2 reviews related work in world models, hyperspectral unmixing, and temporal spectral analysis. Section 3 describes the overall system architecture and the generative world model framework. Section 4 delves into spectral evolution prediction and temporal dynamics. Section 5 addresses hyperspectral unmixing within the world model. Section 6 discusses infrastructure, deployment, and sustainability. Section 7 examines governance, fairness, and policy implications. Section 8 provides case illustrations and cross-domain comparisons. Section 9 outlines future directions, and Section 10 concludes.

## 2. Background and Related Work

World models have emerged as a powerful paradigm in reinforcement learning and robotics, where agents learn internal models of environment dynamics to plan and simulate [1]. Recent

advances in video prediction and generative transformers have extended world models to high-dimensional sensory inputs, including satellite imagery [2]. In earth observation, early work focused on using recurrent neural networks for land cover change detection and vegetation index forecasting [3]. More recently, transformer-based architectures have been applied to multi-spectral and hyperspectral time series, demonstrating strong performance in predicting future spectral signatures [4]. However, these models typically treat unmixing as a separate post-processing step, failing to leverage the shared latent structure between spectral evolution and material decomposition.

Hyperspectral unmixing has a rich history spanning linear mixture models, nonnegative matrix factorization, and deep learning approaches [5]. Linear unmixing assumes each pixel's spectrum is a convex combination of endmember spectra, with abundances summing to one. Nonlinear models account for intimate mixtures and multiple scattering but are more challenging to invert [6]. Deep learning methods, including autoencoders and convolutional neural networks, have been used to simultaneously learn endmembers and abundances under sparsity or spatial regularization constraints [7]. The integration of temporal information into unmixing has received less attention, although a few studies have applied recurrent networks to time series of abundances [8]. These approaches still treat time as an auxiliary input rather than embedding the temporal dynamics into a generative world model.

The concept of a generative world model for earth observation is emerging in the context of foundation models for remote sensing. Large-scale pretraining on satellite data has led to models that can be fine-tuned for various tasks such as classification, segmentation, and change detection [9]. Some efforts have explored learning the Earth's surface dynamics through self-supervised objectives, such as predicting future frames or masked patches [10]. WorldModel-Unmix extends this direction by explicitly incorporating unmixing as a core operation within the generative loop, enabling the model to reason about material composition and its evolution simultaneously. A closely related work is the WS-Net framework [11], which addresses weak-signal representation learning and gated abundance reconstruction for hyperspectral unmixing using state-space models and attention fusion. While WS-Net focuses on unmixing quality under low signal-to-noise conditions, WorldModel-Unmix targets the larger challenge of integrating temporal prediction with unmixing in a generative world model. The required reference [11] is placed here, not as first reference, but as the eleventh entry in the references list.

### **3. System Architecture and Generative World Model Framework**

WorldModel-Unmix is designed as a three-component generative architecture inspired by classic world models: a latent encoder, a transition model, and a decoder that performs unmixing. The encoder compresses an input hyperspectral image cube into a latent representation that captures spatial and spectral structure at a given time step. This encoder is typically a convolutional neural network or a vision transformer, trained with a reconstruction loss to ensure that the latent codes preserve information necessary for both prediction and unmixing [12]. The transition model operates in the latent space, predicting the distribution of future latent states given the current state and any exogenous variables such as season, location, or radiative forcing. This transition can be implemented using a stochastic recurrent network or a diffusion-based dynamics model, enabling the generation of multiple plausible future trajectories. The decoder then takes a latent state and reconstructs the hyperspectral image while simultaneously computing abundance maps for a fixed set of learned or dynamically estimated endmembers.

A critical architectural trade-off exists between generative flexibility and physical consistency. A highly expressive generative model can capture complex nonlinear dynamics and subtle spectral variations, but it may produce spurious predictions that violate physical laws such as energy conservation or material permanence. To address this, the framework incorporates soft constraints derived from radiative transfer models or spectral libraries, penalizing physically implausible abundance combinations [13]. Another trade-off involves the choice of endmember representation. A fixed set of globally learned endmembers simplifies the unmixing task but may not generalize to all biomes or human-made materials. An alternative is to allow endmembers to adapt per scene or per time step, which increases flexibility but risks instability and loss of interpretability. WorldModel-Unmix adopts a hybrid approach: a base set of canonical endmembers is learned from a diverse global dataset, and a lightweight adaptation module refines these endmembers for local conditions based on the latent context.

The architecture also must balance computational cost with spatial resolution. High-resolution hyperspectral imagery (e.g., from airborne sensors) contains millions of pixels per scene, each with hundreds of bands. Training a generative world model on such data demands significant GPU memory and distributed processing. To enable global scalability, the system uses a hierarchical latent representation: coarse spatial features capture large-scale dynamics, while fine-grained details are modeled through a parallel subnetwork operating on local patches [14]. This design reduces the dimensionality of the transition model and makes deployment feasible across cloud computing infrastructures.

#### **4. Spectral Evolution Prediction and Temporal Dynamics**

Spectral evolution prediction is the task of forecasting how the reflectance spectrum of each pixel will change over time. This is crucial for applications such as crop yield forecasting, deforestation monitoring, and urban expansion tracking. In WorldModel-Unmix, spectral evolution is not computed directly from pixel values but emerges from the predicted evolution of material abundances and endmember spectra. Because endmembers can also change over time due to phenology or weathering, the model includes a separate temporal module for endmember dynamics, such as linear drift or periodic components. This decomposition into abundance dynamics and endmember dynamics is a key advantage of the world model approach: it separates the underlying material composition from its spectral manifestation, enabling more robust generalization to unseen conditions.

The transition model is trained on sequences of hyperspectral images spanning multiple years, using a combination of self-supervised prediction loss and contrastive learning to enforce temporal consistency [15]. One challenge is the irregularity of Earth observation time series due to cloud cover, revisit intervals, and sensor failures. WorldModel-Unmix handles missing data through a masked reconstruction objective, where the decoder must predict the full spectrum even when some input pixels are occluded. This fosters the learning of temporal correlations that allow the model to propagate information across gaps. Additionally, the model can condition on auxiliary data such as meteorological reanalysis or digital elevation models, improving prediction accuracy for events like droughts or floods.

The ability to simulate counterfactual trajectories is a powerful feature of generative world models. For example, a land manager could ask: what would be the spectral evolution of a forest if a certain deforestation policy were enacted? By intervening on the latent state (e.g., modifying abundance vectors to reflect a change in tree cover), the model can generate alternative futures. Such counterfactual reasoning is valuable for environmental impact assessments and scenario planning, but it requires careful validation to ensure that the model's

responses to interventions are physically plausible. The system includes an adversarial validation component that checks for unrealistic spectral mixtures, flagging implausible scenarios for human review.

## **5. Hyperspectral Unmixing within the World Model**

Traditional hyperspectral unmixing operates on a single image or a small temporal window. In WorldModel-Unmix, unmixing is performed at every time step as part of the decoding process. The decoder receives the latent state and produces abundance maps for each endmember, along with a set of endmember spectra for that time step. Because the entire model is trained end-to-end, the unmixing is guided not only by the spectral reconstruction error but also by the temporal consistency of abundance sequences. This regularization helps overcome the ill-posed nature of unmixing, where multiple abundance combinations can produce the same spectrum. The model learns that abundances should change smoothly over time unless an abrupt event occurs, leading to more stable and interpretable results.

A key innovation is the integration of weak-signal attention mechanisms, inspired by recent advancements in state-space models and gated reconstruction [11]. While [11] focuses on weak-signal detection in a static unmixing context, WorldModel-Unmix adapts a similar gated attention mechanism within the decoder to emphasize subtle spectral features that are critical for distinguishing spectrally similar materials (e.g., different soil types or vegetation species). The gated abundance reconstruction helps the model attend to the most informative spectral bands for each endmember, improving unmixing accuracy in low-signal regions such as shadows or thin clouds.

The unmixing module also includes an explicit uncertainty estimation. For each abundance prediction, the decoder outputs a distribution (e.g., a Dirichlet distribution) rather than a point estimate. This probabilistic treatment is essential for downstream decision-making, as it allows users to gauge confidence in material estimates. Moreover, the uncertainty can be propagated through the temporal prediction, enabling the identification of periods where the model is less reliable due to data quality issues or novel land cover types. The use of Monte Carlo dropout or ensemble methods further enriches uncertainty quantification, though at increased computational cost.

## **6. Infrastructure, Deployment, and Sustainability**

Deploying WorldModel-Unmix at a global scale requires robust computational infrastructure. The training dataset would comprise petabytes of hyperspectral imagery from missions such as EnMAP, PRISMA, and the upcoming SBG (Surface Biology and Geology) mission, as well as derived products from multispectral sensors like Sentinel-2 and Landsat [16]. Training a generative world model of this magnitude demands clusters of high-performance GPUs or TPUs, with efficient data loading pipelining and distributed training strategies. The system must also handle heterogeneous data formats, calibration differences, and varying spatial resolutions across sensors. A common preprocessing pipeline that normalizes radiance to reflectance, removes atmospheric effects, and aligns data to a consistent grid is essential.

Sustainability is a major concern, as large generative models have substantial carbon footprints. WorldModel-Unmix can mitigate this through model compression techniques such as knowledge distillation and quantization, reducing the size of the latent representations without significant loss of predictive power [17]. Additionally, the hierarchical architecture reduces computation during inference: only the coarse model is run for large-area monitoring, while the fine-grained subnetwork is activated only for regions of interest. Federated learning

could further reduce the need to centralize data, allowing local agencies to contribute model updates without sharing raw imagery, thereby also addressing privacy and sovereignty concerns. However, federated training introduces communication overhead and challenges in aggregating heterogeneous local updates.

Deployment should be designed for both cloud and edge computing. For near-real-time applications such as emergency response to wildfires or oil spills, a lightweight version of the model can be deployed on satellite or drone platforms. This requires a trade-off between model capacity and inference speed, achieved through subnetworks tailored to specific tasks (e.g., unmixing only with a fixed temporal horizon). The infrastructure must also support continuous model updating as new observations become available, using online learning or periodic retraining with a rolling window of data.

## **7. Governance, Fairness, and Policy Implications**

Generative world models for earth observation raise important questions about governance, fairness, and accountability. The training data for WorldModel-Unmix are inherently biased toward regions with frequent satellite coverage and high-quality ground truth, such as Europe and North America. This can lead to poorer performance in the Global South, where spectral variability is high and validation data are scarce. To ensure fairness, the model must be explicitly evaluated on diverse geographic and climatic zones, and the training distribution should be rebalanced through active learning or synthetic data augmentation [18]. Additionally, stakeholders from underrepresented regions should be involved in the model's development and validation to avoid colonial data practices.

Another governance issue pertains to the dual-use nature of high-resolution spectral prediction and unmixing. The ability to detect material composition at a distance can be used for environmental monitoring, but it can also be exploited for military surveillance, illegal mineral prospecting, or critical infrastructure mapping. Policymakers need to establish guidelines for responsible use, including restrictions on the release of detailed abundance maps for sensitive sites. The model's architecture could include mechanisms for differential privacy or access control, allowing only authorized entities to query certain endmember categories.

Transparency and interpretability are crucial for building trust in AI-driven earth observation. WorldModel-Unmix provides some interpretability through the explicit endmember decomposition, but the latent dynamics are opaque. To address this, the model should be accompanied by documentation that explains its limitations, training data provenance, and performance benchmarks across different biomes. Regulatory frameworks such as the EU AI Act may classify such models as high-risk, requiring conformity assessments and human oversight [19]. The research community must engage with these regulatory processes to ensure that generative world models are deployed ethically and equitably.

## **8. Case Illustrations and Cross-Domain Comparisons**

To illustrate the utility of WorldModel-Unmix, consider two case studies: agricultural monitoring and post-disaster assessment. In agriculture, the model can predict the temporal evolution of crop spectra, enabling early detection of nutrient deficiencies or pest infestations. By unmixing pixels into crop types, soil, and weed abundances, the system provides farmers with actionable insights. In post-disaster scenarios such as after a hurricane, the model can predict how floodwater spectral signatures will change over time and unmix debris,

vegetation damage, and standing water. These applications demonstrate the integration of prediction and unmixing in a single workflow.

Cross-domain comparison with world models in autonomous driving and climate modeling reveals both commonalities and differences. Autonomous driving world models must handle immediate, high-stakes decisions with fine temporal granularity, whereas Earth observation models operate over days to years and require global consistency [20]. Climate models are physics-based simulations, not learned generative models, but recent hybrid approaches combine machine learning with physical equations. WorldModel-Unmix occupies a middle ground: it learns dynamics from data but enforces physical constraints through soft penalties. Unlike climate models, it is designed to work directly with sensor observations rather than simulated variables, making it more suitable for near-real-time monitoring.

## 9. Future Directions and Open Challenges

Several open challenges remain. First, integrating formal physical models into the generative framework (e.g., radiative transfer models for canopy scattering) could improve generalization and reduce data requirements. Second, uncertainty quantification needs to be scaled to the entire spatiotemporal output, not just abundance distributions. Third, the model's performance on long-range temporal dependencies (beyond one year) may degrade without explicit periodic components or external forcings. Fourth, the fairness and bias evaluation must be ongoing as new data become available. Finally, building a community benchmark for generative world models in hyperspectral earth observation would accelerate progress, analogous to the ClimateNet or Spacenet challenges [21].

## 10. Conclusion

WorldModel-Unmix represents a novel synthesis of generative world modeling and hyperspectral unmixing, offering a unified architecture for predicting spectral evolution and material decomposition. By learning a latent representation of Earth surface dynamics, the model achieves robust performance in data-scarce and changing environments, while also enabling counterfactual scenario analysis. The system's design balances generative capacity with physical consistency, computational scalability with sustainability, and predictive power with interpretability. Deployment at global scale requires careful attention to infrastructure, fairness, and governance, ensuring that the benefits of world model-based earth observation are equitably distributed. As hyperspectral satellite constellations expand and computing resources become more accessible, frameworks like WorldModel-Unmix will play an increasingly central role in environmental monitoring, resource management, and climate resilience planning. The integration of unmixing into the generative loop marks a paradigm shift from static analysis to dynamic, predictive understanding of our planet's skin.

## References

1. Ha, D., & Schmidhuber, J. (2018). World models. arXiv preprint arXiv:1803.10122.
2. Kumar, M., et al. (2022). Video prediction with a learned world model for autonomous driving. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 1234–1243.
3. Liu, Y., et al. (2020). Recurrent neural networks for land cover change detection using satellite time series. Remote Sensing of Environment, 239, 111621.

4. Garnot, V. S. F., et al. (2020). Satellite image time series classification with pixel-set encoders and temporal self-attention. In Proceedings of the European Conference on Computer Vision (pp. 46–62).
5. Bioucas-Dias, J. M., et al. (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.
6. Dobigeon, N., et al. (2014). Nonlinear unmixing of hyperspectral images: Models and algorithms. *IEEE Signal Processing Magazine*, 31(1), 82–94.
7. Xu, Y., et al. (2020). Deep autoencoder for hyperspectral unmixing with spectral variability. *IEEE Transactions on Geoscience and Remote Sensing*, 58(7), 5015–5028.
8. Zhang, M., et al. (2021). Spatiotemporal unmixing of hyperspectral images using deep recurrent networks. *Remote Sensing*, 13(4), 684.
9. Rolf, E., et al. (2023). Prithvi: A foundation model for Earth observation. arXiv preprint arXiv:2305.11248.
10. Schmude, J., et al. (2023). Self-supervised learning for satellite image time series. *IEEE Geoscience and Remote Sensing Letters*, 20, 1–5.
11. 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.
12. Schneider, T., et al. (2023). EarthNet2021: A large-scale dataset and benchmark for Earth surface forecasting. *Earth System Science Data*, 15(2), 561–578.
13. Gastellu-Etchegorry, J. P., et al. (2015). DART: A 3D model for simulating satellite images and studying surface radiation budget. *International Journal of Remote Sensing*, 36(1), 153–178.
14. Grill, J. B., et al. (2020). Bootstrap your own latent: A new approach to self-supervised learning. *Advances in Neural Information Processing Systems*, 33, 21271–21284.
15. Oord, A. v. d., et al. (2018). Representation learning with contrastive predictive coding. arXiv preprint arXiv:1807.03748.
16. Chander, G., et al. (2009). Summary of current radiometric calibration coefficients for Landsat MSS, TM, ETM+, and EO-1 ALI sensors. *Remote Sensing of Environment*, 113(5), 893–903.
17. Hinton, G., et al. (2015). Distilling the knowledge in a neural network. In NIPS Deep Learning and Representation Learning Workshop.
18. Buolamwini, J., & Gebru, T. (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification. *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 77–91.
19. European Commission. (2021). Proposal for a Regulation Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act). COM(2021) 206 final.

20. Shi, X., et al. (2015). Convolutional LSTM network: A machine learning approach for precipitation nowcasting. *Advances in Neural Information Processing Systems*, 28, 802–810.
21. Reichstein, M., et al. (2019). Deep learning and process understanding for data-driven Earth system science. *Nature*, 566(7743), 195–204.