

CROPLAND MAPPING USING GEOSPATIAL EMBEDDINGS

Ivan Zvonkov¹ Gabriel Tseng^{2,3} Inbal Becker-Reshef¹ Hannah Kerner⁴

¹ University of Maryland, College Park, ² Mila – Quebec AI Institute,

³ McGill University, ⁴ Arizona State University

ABSTRACT

Accurate and up-to-date land cover maps are essential for understanding land use change, a key driver of climate change. Geospatial embeddings offer a more efficient and accessible way to map landscape features, yet their use in real-world mapping applications remains underexplored. In this work, we evaluated the utility of geospatial embeddings for cropland mapping in Togo. We produced cropland maps using embeddings from Presto and AlphaEarth. Our findings show that geospatial embeddings can simplify workflows, achieve high-accuracy cropland classification and ultimately support better assessments of land use change and its climate impacts.

1 INTRODUCTION & RELATED WORK

Human-driven land transformation is the most substantial human alteration of the Earth (Vitousek et al., 1997). Researchers have used Earth observation data to better understand land use change (e.g., measuring global cropland expansion using the Landsat satellite data archive (Potapov et al., 2022)). The amount of data that needs to be fetched and processed presents a major bottleneck for land use and land change mapping efforts.

Geospatial embeddings provide a fast and scalable way to turn Earth observation data into actionable maps. They are created by first training a model with self-supervised learning and then running large-scale inference over a region of interest (e.g., (Feng et al., 2025)). Models that generate geospatial embeddings are often called geospatial foundation models (GeoFMs). While many GeoFMs have been released (Astruc et al., 2024; Fuller et al., 2023; Xiong et al., 2024; Nedungadi et al., 2024; Cong et al., 2022; Reed et al., 2022; Wang et al., 2024; Jakubik et al., 2023), their use in real-world applications remains limited (Ghamisi et al., 2025).

In this paper, we explored the use case of cropland mapping using geospatial embeddings. We generated cropland maps in Togo using embeddings from Presto (Tseng et al., 2023) and AlphaEarth (Brown et al., 2025). We used Presto because it has relatively low compute requirements while maintaining strong performance across Earth observation tasks. We used AlphaEarth (Brown et al., 2025) embeddings because they have already been generated globally. To evaluate utility, we analyzed the quality of the generated cropland maps and compared them to existing products. By lowering barriers to generating accurate cropland maps, geospatial embeddings can support more timely monitoring of land use change and its impact on the climate.

2 METHODOLOGY

2.1 EMBEDDING GENERATION

We generated embeddings for all of Togo (56,785 km²) for the time frame of March 2019 - March 2020. Embedding generation took 16 hours and cost \$313.40. The final Togo asset size was 128.8 GB. The Google Earth Engine (GEE) asset and all relevant scripts are available at: https://github.com/ivanzvonkov/ML4RS_2026_Tutorial.

The embedding generation pipeline involved two steps: model deployment and inference. For model deployment, we wrote a Google Colab notebook to package Presto into a TorchServe Docker

container and deploy the container to Google Cloud’s Vertex AI. For inference, we created a GEE script to obtain Earth observation data for all of Togo and make predictions from the data using the `ee.Model.fromVertexAi` function. We used the Presto encoder to compress location information, optical imagery (Sentinel-2), radar imagery (Sentinel-1), climatology data (ERA5), and elevation data (SRTM) over the course of a year. Each generated Presto embedding contained 128 features representing a single 10m pixel on Earth.

To verify that our generated embeddings contained meaningful information, we clustered the embeddings and visually compared them to the WorldCover land cover map (Zanaga et al., 2022), which has been shown to have accurate cropland classification in Togo (Kerner et al., 2024). We performed clustering using the `ee.Clusterer.wekaKMeans` algorithm available in GEE with 7 classes. We sampled 1000 points from our generated embeddings to train the clustering algorithm.

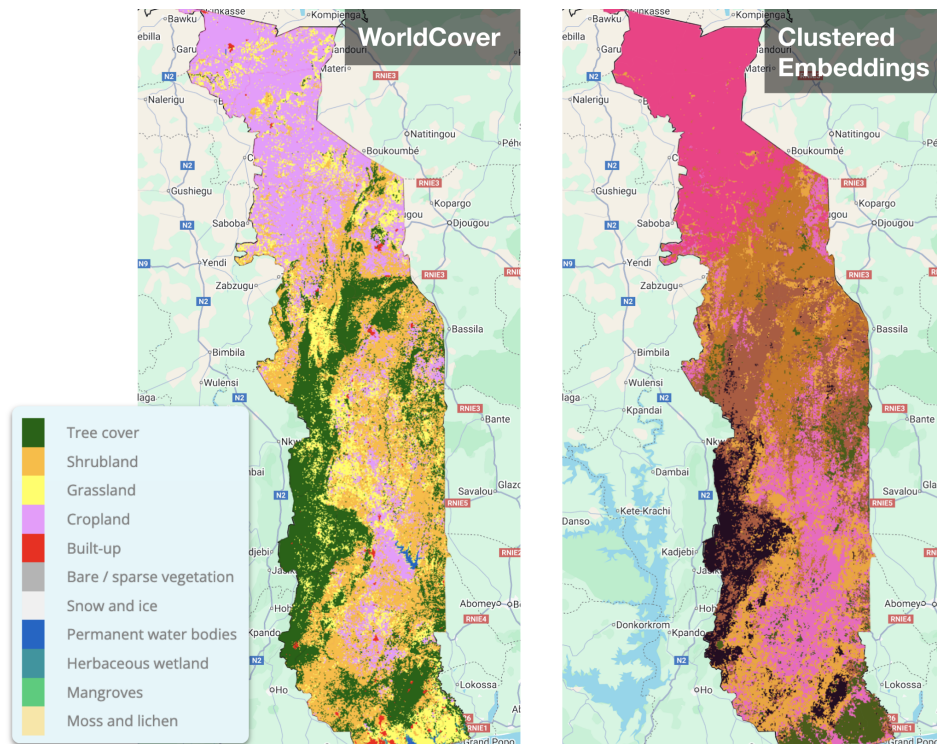


Figure 1: Comparison of WorldCover land cover map and clustered embeddings in Togo. We used the default WorldCover color scheme. The cluster colors were randomly generated.

Figure 1 showed that the Presto embedding clusters were visually aligned with the land cover map. In particular, the green embedding cluster correlated closely with built-up areas in the land cover map. This gave us confirmation that the embeddings captured meaningful information about the land cover in Togo and could be used for further mapping and analysis.

2.2 MAPPING USING EMBEDDINGS

Our embedding approach allowed us to train a classifier using **8 lines of code** (see the GEE script below). When training a classifier on top of the embeddings, we used labeled points from the Togo cropland classification task in CropHarvest (Tseng et al., 2021). We evaluated the classifier’s outputs using labels collected in March 2020 as part of a Rapid Response crop mapping effort in Togo (Kerner et al., 2020). We used the default GEE Random Forest classifier initialized with 100 trees. We trained the classifier on a training dataset consisting of Togo labels paired with embeddings. The initial outputs contained extensive crop predictions, so we applied a probability threshold of 0.7 to reduce false positives. The resulting cropland map was generated in less than 5 seconds. The entire script is included below:

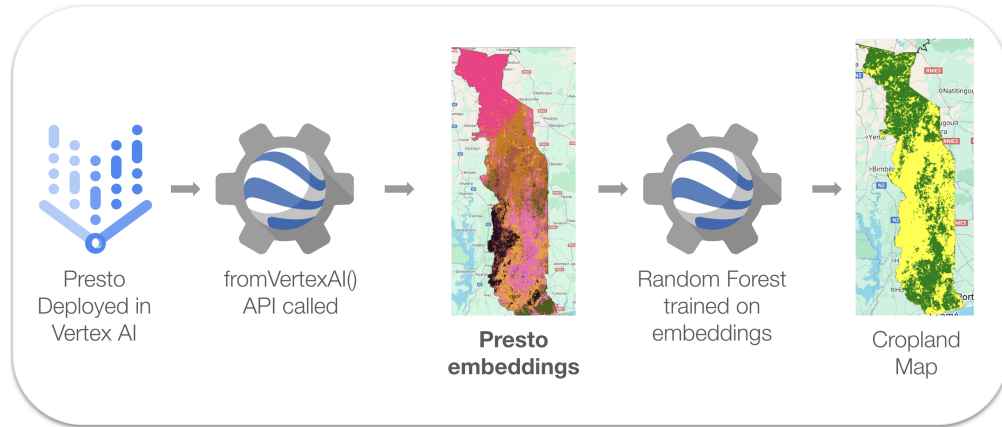


Figure 2: Steps involved in generating embeddings and mapping cropland.

```
// 1. Load embeddings for region of interest
var roi = ee.FeatureCollection("FAO/GAUL/2015/level2").filter("ADM0_NAME
=='Togo'");
var embeddings = ee.Image("users/<omitted for review>/Togo/
Presto_embeddings_v2025_06_19")

// 2. Load Togo points
var points = ee.FeatureCollection("users/<omitted for review>/Togo/
points_2019")
var trainPoints = points.filter(ee.Filter.eq("subset", "training"))

// 3. Create training dataset (training points + embeddings)
var trainSet = embeddings.sampleRegions(trainPoints, ["is_crop"], 10)

// 4. Train a classifier (100 trees)
var model = ee.Classifier.smileRandomForest(100).setOutputMode('
probability').train(trainSet, 'is_crop', embeddings.bandNames());

// 5. Classify embeddings using trained model
var croplandPreds = embeddings.classify(model).clip(roi)
var croplandMap = croplandPreds.gte(0.7).rename("map_crop")
```

Our methodology is illustrated in Figure 2. For comparison, we applied the same cropland mapping process to the AlphaEarth embeddings (Brown et al., 2025) (comparison of Presto and AlphaEarth embeddings is shown in Table 1).

	Presto Embeddings	AlphaEarth Embeddings
Data Sources	Sentinel-1, Sentinel-2, ERA5, SRTM	Optical, radar, LiDAR, and other sources
Time frame	March 2019 – March 2020	January 2021 – January 2022 [earlier than 2021 not available]
Embedding Size	128 values	64 values
Scale	10 m	10m
Data Type	uint16	double
Bytes per pixel	256	512

Table 1: Comparison of Presto and AlphaEarth embeddings.

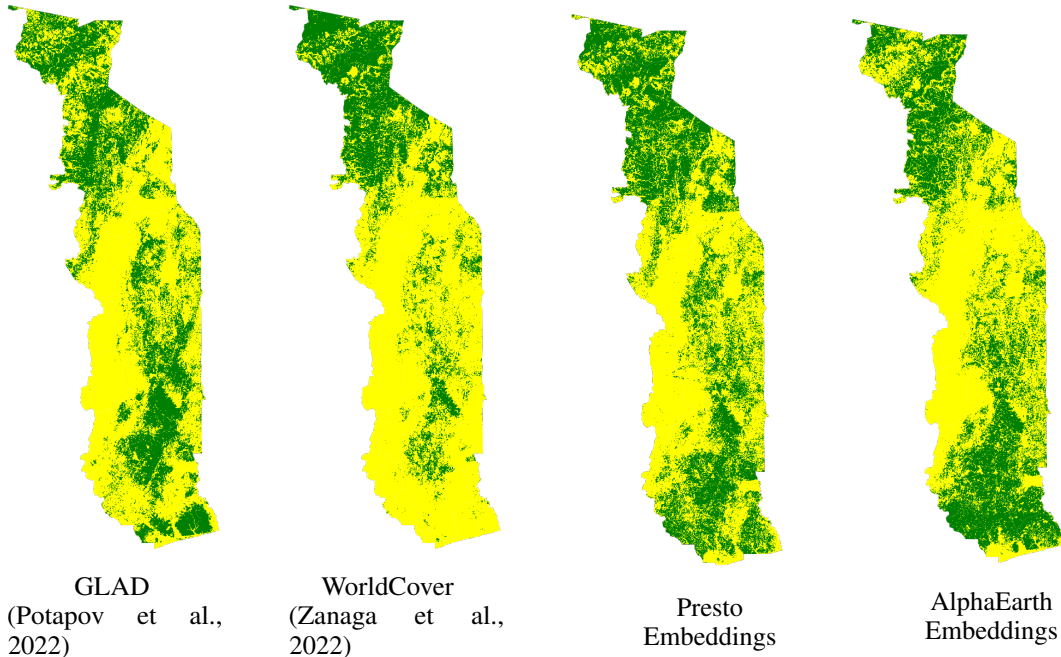


Figure 3: Comparison of cropland maps in Togo. Green is cropland and yellow is not cropland.

3 RESULTS & CONCLUSION

We conducted a visual assessment on the generated cropland maps and compared them to two existing cropland maps (Figure 3). For the existing cropland maps, we chose GLAD (Potapov et al., 2022) and WorldCover (Zanaga et al., 2022) because they have high cropland accuracy in Togo (Kerner et al., 2024). We found that Presto and AlphaEarth embedding-based maps have similar cropland proportion and distribution. We also noted that Presto and AlphaEarth embedding-based maps classify more pixels as cropland than WorldCover. The appendix provides additional qualitative analysis.

To get a more comprehensive understanding of the map quality, we conducted an accuracy assessment using the Togo test set from CropHarvest (Tseng et al., 2021). It’s important to note that both Presto and AlphaEarth embeddings were benchmarked against data in the Togo test set. The Presto embeddings-based cropland map achieved the highest overall accuracy, producer’s accuracy, and F1 score (Table 2).

	GLAD (Potapov et al., 2022)	WorldCover (Zanaga et al., 2022)	Presto embeddings	AlphaEarth embeddings
Overall Accuracy	0.859	0.880	0.897	0.859
User’s Accuracy	0.821	0.892	0.833	0.745
Producer’s Accuracy	0.627	0.647	0.784	0.745
F1 Score	0.711	0.750	0.808	0.745

Table 2: Cropland metrics for Togo.

Conclusion We showed that using embeddings for cropland mapping is simple, quick, and effective. While our experiments are limited to Togo, they serve as a case study to illustrate the methodology. Future work will extend this evaluation across additional regions and applications. Our results highlight the potential for using embeddings for more effective mapping and demonstrate the growing potential of embeddings for understanding land use change and its impact on the climate.

REFERENCES

- Guillaume Astruc, Nicolas Gonthier, Clement Mallet, and Loic Landrieu. AnySat: An Earth observation model for any resolutions, scales, and modalities. *arXiv preprint arXiv:2412.14123*, 2024.
- Christopher Brown, Michal Kazmierski, Valerie Pasquarella, William Rucklidge, Masha Samsikova, Chenhui Zhang, Evan Shelhamer, Estefania Lahera, Olivia Wiles, Simon Ilyushchenko, Noel Gorelick, Lihui Zhang, Sophia Alj, Emily Schechter, Sean Askay, Oliver Guinan, Rebecca Moore, Alexis Boukouvalas, and Pushmeet Kohli. Alphaearth foundations: An embedding field model for accurate and efficient global mapping from sparse label data, 07 2025.
- Yezen Cong, Samar Khanna, Chenlin Meng, Patrick Liu, Erik Rozi, Yutong He, Marshall Burke, David B. Lobell, and Stefano Ermon. SatMAE: Pre-training transformers for temporal and multi-spectral satellite imagery. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *NeurIPS*, 2022. URL <https://openreview.net/forum?id=WBhqzpf6KYH>.
- Zhengpeng Feng, Clement Atzberger, Sadiq Jaffer, Jovana Knezevic, Silja Sormunen, Robin Young, Madeline C Lisaius, Markus Immitzer, David A Coomes, Anil Madhavapeddy, et al. Tessera: Temporal embeddings of surface spectra for earth representation and analysis. *arXiv preprint arXiv:2506.20380*, 2025.
- Anthony Fuller, Koreen Millard, and James R Green. CROMA: Remote sensing representations with contrastive radar-optical masked autoencoders. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=eqI5WgGvY>.
- Pedram Ghamisi, Weikang Yu, Xiaokang Zhang, Aldino Rizaldy, Jian Wang, Chufeng Zhou, Richard Gloaguen, and Gustau Camps-Valls. Geospatial foundation models to enable progress on sustainable development goals. 05 2025. doi: 10.48550/arXiv.2505.24528.
- Johannes Jakubik, Sujit Roy, C. E. Phillips, Paolo Fraccaro, Denys Godwin, Bianca Zadrozny, Daniela Szwarcman, Carlos Gomes, Gabby Nyirjesy, Blair Edwards, Daiki Kimura, Naomi Simumba, Linsong Chu, S. Karthik Mukkavilli, Devyani Lambhate, Kamal Das, Ranjini Bangalore, Dario Oliveira, Michal Muszynski, Kumar Ankur, Muthukumaran Ramasubramanian, Iksha Gurung, Sam Khallaghi, Hanxi (Steve) Li, Michael Cecil, Maryam Ahmadi, Fatemeh Kordi, Hamed Alemohammad, Manil Maskey, Raghu Ganti, Kommy Weldemariam, and Rahul Ramachandran. Foundation Models for Generalist Geospatial Artificial Intelligence. *Preprint Available on arxiv:2310.18660*, October 2023.
- Hannah Kerner, Gabriel Tseng, Inbal Becker-Reshef, Catherine Nakalembe, Brian Barker, Blake Munshell, Madhava Paliyam, and Mehdi Hosseini. Rapid response crop maps in data sparse regions. In *ACM SIGKDD Conference on Data Mining and Knowledge Discovery Workshops*, 2020.
- Hannah Kerner, Catherine Nakalembe, Adam Yang, Ivan Zvonkov, Ryan McWeeny, Gabriel Tseng, and Inbal Becker-Reshef. How accurate are existing land cover maps for agriculture in sub-saharan africa? *Scientific Data*, 11(1):486, 2024.
- Vishal Nedungadi, Ankit Kariryaa, Stefan Oehmcke, Serge Belongie, Christian Igel, and Nico Lang. Mmearth: Exploring multi-modal pretext tasks for geospatial representation learning. In *European Conference on Computer Vision*, pp. 164–182. Springer, 2024.
- Peter Potapov, Svetlana Turubanova, Matthew C. Hansen, Alexander Tyukavina, Victoria Zalles, Amrit Khan, Xiao-Peng Song, Andrew Pickens, Qi Shen, and Juan Cortez. Global maps of cropland extent and change show accelerated cropland expansion in the twenty-first century. *Nature Food*, 3:19–28, 2022. doi: 10.1038/s43016-021-00429-z. URL <https://doi.org/10.1038/s43016-021-00429-z>.
- Colorado J Reed, Ritwik Gupta, Shufan Li, Sarah Brockman, Christopher Funk, Brian Clipp, Salvatore Candido, Matt Uyttendaele, and Trevor Darrell. Scale-mae: A scale-aware masked autoencoder for multiscale geospatial representation learning. *arXiv preprint arXiv:2212.14532*, 2022.

Gabriel Tseng, Ivan Zvonkov, Catherine Lilian Nakalembe, and Hannah Kerner. Cropharvest: A global dataset for crop-type classification. In *NeurIPS, Datasets and Benchmarks Track*, 2021. URL <https://openreview.net/forum?id=JtjzUXPEaCu>.

Gabriel Tseng, Ruben Cartuyvels, Ivan Zvonkov, Mirali Purohit, David Rolnick, and Hannah R Kerner. Lightweight, pre-trained transformers for remote sensing timeseries. In *NeurIPS 2023 Workshop on Tackling Climate Change with Machine Learning*, 2023. URL <https://www.climatechange.ai/papers/neurips2023/58>.

Peter M. Vitousek, Harold A. Mooney, Jane Lubchenco, and Jerry M. Melillo. Human domination of earth’s ecosystems. *Science*, 277(5325):494–499, 1997. doi: 10.1126/science.277.5325.494. URL <https://www.science.org/doi/abs/10.1126/science.277.5325.494>.

Yi Wang, Conrad M Albrecht, and Xiao Xiang Zhu. Multi-label guided soft contrastive learning for efficient earth observation pretraining, 2024.

Zhitong Xiong, Yi Wang, Fahong Zhang, Adam J Stewart, Joëlle Hanna, Damian Borth, Ioannis Papoutsis, Bertrand Le Saux, Gustau Camps-Valls, and Xiao Xiang Zhu. Neural plasticity-inspired foundation model for observing the Earth crossing modalities. *arXiv preprint arXiv:2403.15356*, 2024.

Daniele Zanaga, Ruben Van De Kerchove, D Daems, W De Keersmaecker, Carsten Brockmann, G Kirches, Jan Wevers, O Cartus, M Santoro, S Fritz, et al. ESA WorldCover 10 m 2021 v200, 2022.

A APPENDIX

A.1 ADDITIONAL QUALITATIVE ANALYSIS

We conducted further investigation into the South of Togo where all maps have differences. We zoomed into a window centered on coordinate: (0.84445, 6.46924) and compared a high resolution satellite image (Figures 4, 7) to all the cropland maps (Figures 5, 6, 8, 9). In this window we found that all maps with the exception of GLAD correctly identify built-up areas as non-crop. Embedding-based maps appear to classify more cropland (green) correctly. AlphaEarth embedding-based map over predicts crops in some shrub/tree areas, while the Presto embedding-based map visually accurately identifies many of the tree areas as non-crop.




Figure 4: High resolution satellite image zoomed in on coordinate: (0.84445, 6.46924).




Figure 5: GLAD cropland map zoomed in on coordinate: (0.84445, 6.46924).




Figure 6: WorldCover cropland map zoomed in on coordinate: (0.84445, 6.46924).




Figure 7: High resolution satellite image zoomed in on coordinate: (0.84445, 6.46924).




Figure 8: Presto cropland map zoomed in on coordinate: (0.84445, 6.46924).




Figure 9: AlphaEarth cropland map zoomed in on coordinate: (0.84445, 6.46924).