MONITORING ILLICIT RARE EARTH MINING IN MYANMAR VIA SELF-SUPERVISED LEARNING

Ollie Ballinger
Centre for Advanced Spatial Analysis
University College London
90 Tottenham Court Road, W1T 4TJ
oballinger@ucl.ac.uk

ABSTRACT

Heavy Rare Earth Elements (HREEs) are critical for the production of most electronic devices. Rapidly increasing demand for these minerals has led to a proliferation of highly polluting makeshift HREE extraction in Myanmar. Monitoring the spread of these mines is important for the preservation of human health and the environment. This paper utilizes a geospatial foundation model pre-trained using self-supervised learning to detect hundreds of rare earth mines using a single template example. This is achieved through the development of a novel method for embedding similarity search– Cosine Contrast– which leverages both positive and negative templates to yield more relevant results.

1 INTRODUCTION

Heavy Rare Earth Elements (HREEs) are essential components in the production of electric vehicles, military hardware, and consumer electronics (U.S. Geological Survey, 2014). Demand for these metals is projected to increase fivefold by 2040 (Canada, 2025). Recently, China's imports of HREEs from Myanmar have risen from $1.5 million in 2014 to $1.4 billion in 2023 (Witness, 2022). An investigation by Global Witness found extensive illicit production and trade of these minerals in rebel-controlled areas along Myanmar's northeastern border with China (Witness, 2024). The unregulated mining of rare earth elements poses significant risks to human health and the environment.

Several studies have used supervised learning to detect similar Artisanal Small-scale Mining (ASM) operations, often in the Amazon or Sub-Saharan Africa (Ballinger, 2023; al, 2016; Kimijima et al., 2021; Nava et al., 2022; Nyamekye et al., 2021; Simionato et al., 2021). For example, (Owusu-Nimo et al., 2018) developed a deep convolutional neural network applied to Sentinel-2 satellite imagery in Ghana, achieving high accuracy. Their approach demonstrated the ability to monitor ASM, including quantifying illegal mining-related deforestation within protected forests. However, supervised learning often requires hundreds if not thousands of labeled examples in order to perform.

Recent progress has been made in the application of Self-Supervised Learning (SSL) techniques in earth observation (Jakubik et al., 2023). These methods require no further labels once pre-trained, and generate numerical representations of the features contained in patches of satellite imagery, called embeddings. This paper explores the use of a pre-trained geospatial foundation model to

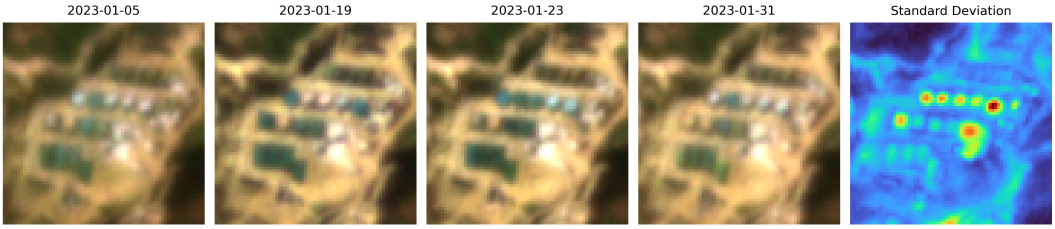

Figure 1: HREE precipitation reaction in 3m PlanetScope imagery

detect HREE mining in Myanmar, and proposes Cosine Contrast as a novel method for the retrieval of semantically similar embeddings. The proposed method achieves high accuracy (AUC=0.95) using a single example of HREE mine to identify hundreds of others in the area.

2 METHODOLOGY

2.1 DATA

The main source of data used in this study is PlanetScope satellite imagery, which has a spatial resolution of 3 meters per pixel and a daily repeat cycle. One main collection on 2022-03-07 was used for inference, covering an 1,800 km^2 area north of Pang War, on the China-Myanmar border. All eight spectral bands were employed. Because HREE precipitation pools are approximately 10-15 meters in diameter, 10 meter Sentinel-2 imagery would be too coarse to identify individual pools. Figure 1 displays PlanetScope imagery showing the HREE precipitation reaction occurring at a mine near Pang War, Myanmar. The pools cycle from blue to white over the course of 5-6 days. Pixel standard deviation over time reveals which pools are in use.

For validation, the coordinates of 2,792 HREE precipitation pools were used. These labels were generated manually by Global Witness in March of 2022, as part of their investigation (Witness, 2022).

2.2 SELF-SUPERVISED LEARNING

Self-supervised learning for Earth observation leverages unlabeled satellite imagery to pre-train models by learning meaningful representations through tasks such as image reconstruction (Jakubik et al., 2023). These models can then be fine-tuned with limited labeled data for specific applications such as land cover classification, object detection, and change detection (Jakubik et al., 2023). This approach significantly reduces the need for extensive labeled datasets, enabling more scalable and efficient analysis of remote sensing data.

Recently, a number of global "foundation models" have been developed by pre-training on large geospatial datasets using MAE, including Stanford's SatMAE (Cong et al., 2023), NASA and IBM's Prithvi model (Jakubik et al., 2023), the Allen Institute's Satlas model (Bastani et al., 2023), and the Clay Foundation Model (Clay, 2025), which was used herein.

Clay was pre-trained using a Masked Autoencoder on 33.8 Terabytes of satellite imagery from a range of modalities, including PlanetScope imagery (Clay, 2025). This was one of the main reasons Clay was selected for this task, as PlanetScope imagery for Myanmar is freely available and high enough resolution (3m) to identify individual HREE precipitation pools, as shown in Figure 1.

2.3 EMBEDDING EXTRACTION

To extract embeddings, the Planetscope mosaic of the study area was divided into evenly sized 250m-by-250m (84 pixel) patches. This patch size was selected on the basis of the observed size of HREE mines in the area. Of the resulting 28,757 patches, 531 (roughly 2%) contained at least one HREE precipitation pool, with some containing up to 23 pools.

The Clay v1 model outputs 768-dimensional embedding vectors that encode features within the patches. Figure 3a reduces the embeddings to two dimensions using the Uniform Manifold Approximation and Projection (UMAP) (McInnes et al., 2020), with patches containing a HREE mine colored in red and sized according to the number of precipitation pools therein. When plotting the patches in embedding space– even when reduced to two dimensions– semantic differences between these two sets of patches become clear.

2.4 EMBEDDING SEARCH: COSINE CONTRAST

Once the patch embeddings have been generated, these can be used to search for similar embeddings. A standard approach for similarity search between vectors is to compute the cosine similarity (Xia et al., 2015). Given two embeddings–$\mathbf{E}_j$, which was generated from a patch known to contain

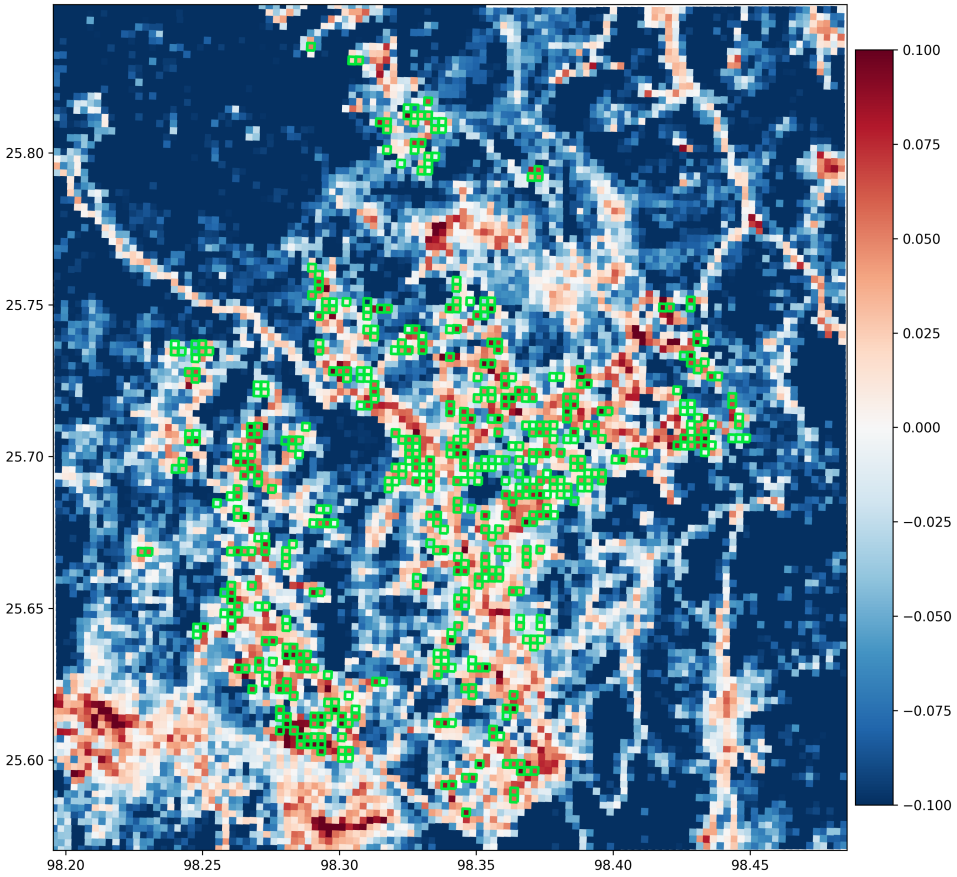

Figure 2: Cosine Contrast ($k = 1$) for 250m patches, HREE mines indicated in green

a HREE mine ($m(\mathbf{E}_j) = 1$), and $\mathbf{E}_i$ for which the value of m is unknown– cosine similarity is calculated as follows:

$$\cos\theta_{i,j} = \frac{\mathbf{E}_i \cdot \mathbf{E}_j}{\|\mathbf{E}_i\|\|\mathbf{E}_j\|} \tag{1}$$

A high value of $\cos\theta_{i,j}$ would suggest that $\mathbf{E}_i$ also contains a HREE mine. In practice, we may have a number (k) of template embeddings that comprise a set P of positive examples:

$$P = \{\mathbf{E}_1^P, \mathbf{E}_2^P, \ldots, \mathbf{E}_k^P \mid m(\mathbf{E}_j^P) = 1 \, \forall j \in \{1, 2, \ldots, k\}\}$$

Calculating the average cosine similarity between $\mathbf{E}_i$ and the set of positive templates P yields an average similarity score:

$$\mu_{i,P} = \frac{1}{|P|} \sum_{j \in P} \cos\theta_{i,j} \tag{2}$$

Calculating average similarity using a set of diverse positive examples improves the retrieval of relevant information by introducing variation in irrelevant features (e.g., background or surrounding vegetation). Extending this logic, detection would be further improved by explicitly contrasting $\mathbf{E}_i$ with a set of negative templates N– embeddings generated from patches in which no mine is present:

$$N = \{\mathbf{E}_1^N, \mathbf{E}_2^N, \ldots, \mathbf{E}_k^N \mid m(\mathbf{E}_j^N) = 0 \, \forall j \in \{1, 2, \ldots, k\}\}$$

Following the calculation of the average similarity between $\mathbf{E}_i$ and the positive templates $\mu_{i,P}$ and the negative templates $\mu_{i,N}$, the Cosine Contrast (Δ_i) can be calculated as the normalized difference between the two:

$$\Delta_i = \frac{\mu_{i,P} - \mu_{i,N}}{\mu_{i,P} + \mu_{i,N}} \tag{3}$$

This process yields values that range from -1 to 1, where positive values indicate that the corresponding embedding is more similar to the positive template than the negative template.

3 RESULTS

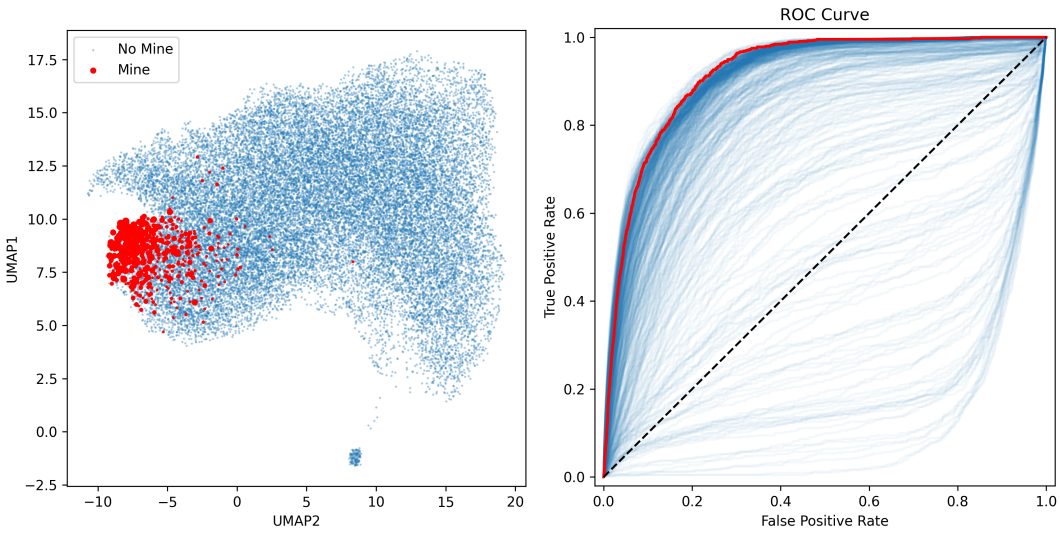

(a) UMAP Projection of Embeddings into two dimensions

(b) ROC Curves for cosine similarity (blue) and Cosine Contrast (red)

Figure 3: Comparison of UMAP projection and ROC curves

Figure 2 shows Cosine Contrast values for patches in the study area, using one positive and one negative template ($k = 1$). Patches containing a HREE mine are indicated in green. There is strong alignment between positive Cosine Contrast values, shown in red, and HREE mine presence. If patches with negative Cosine Contrast values (blue areas) are discarded, 95% of the patches containing mines are retained, while the overall sample size is reduced by 82%. The process could be repeated iteratively with the remaining patches, allowing an end user to implicitly highlight relevant features by select negative examples and further increase the purity of the results.

Figure 3b plots the Receiver Operating Characteristic curves for the task of discriminating between patches that contain an HREE mine and those that don't. The blue ROC curves were generated using the cosine similarity score using each patch containing a mine as the positive template. While this often generates good results, the efficiency of cosine similarity-based embedding search for mine detection is sensitive to the choice of the template example; The average AUC using a single template is 0.86±0.14. When certain templates are used, AUC drops below 0.5, meaning it performs worse than a random classifier.

However, when Cosine Contrast is used, AUC for mine detection reaches 0.95, nearly 10% higher than the average AUC when only a positive template is used. Simply setting a threshold value of Cosine Contrast to classify patches into mine/no mine is 98% accurate, with an F1 score of 0.4.

4 CONCLUSION

The use of supervised learning requires a large number of labeled instances of such mines, which are unlikely to be readily available. The patch embeddings generated by Self-Supervised Learning models enable the efficient detection of similar ground targets using a single template image, though performance can be variable. The use of negative examples to calculate Cosine Contrast further improves accuracy. In the study area, sorting all patches in descending order of Cosine Contrast places 90% of the HREE precipitation pools in the first 10% of the ordered data. This enables efficient wide-area search for HREE precipitation pools.

REFERENCES

Pelletier et al. Assessing the robustness of Random Forests to map land cover with high resolution satellite image time series over large areas. *Remote Sens. Environ.*, 187:156–168, 2016. doi: 10.3390/rs9101065.

Ollie Ballinger. Refinery Identification, 2023. URL https://alan-turing-institute.github.io/DSGFeb2023GW/refineries.html.

Favyen Bastani, Piper Wolters, Ritwik Gupta, Joe Ferdinando, and Aniruddha Kembhavi. SatlasPretrain: A Large-Scale Dataset for Remote Sensing Image Understanding, August 2023. URL http://arxiv.org/abs/2211.15660. arXiv:2211.15660 [cs].

Natural Resources Canada. Critical minerals value chains – Permanent magnets, January 2025. URL https://www.canada.ca/en/campaign/critical-minerals-in-canada/critical-minerals-an-opportunity-for-canada/permanent-magnets.html. Last Modified: 2025-01-09.

Clay. Clay Foundation Model — Clay Foundation Model, 2025. URL https://clay-foundation.github.io/model/index.html#.

Yezhen 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, January 2023. URL http://arxiv.org/abs/2207.08051. arXiv:2207.08051 [cs].

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, 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, November 2023. URL http://arxiv.org/abs/2310.18660. arXiv:2310.18660 [cs].

Satomi Kimijima, Masayuki Sakakibara, Masahiko Nagai, and Nurfitri Abdul Gafur. Time-Series Assessment of Camp-Type Artisanal and Small-Scale Gold Mining Sectors with Large Influxes of Miners Using LANDSAT Imagery. *International Journal of Environmental Research and Public Health*, 18(18):9441, September 2021. doi: 10.3390/ijerph18189441. MAG ID: 3197035829.

Leland McInnes, John Healy, and James Melville. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction, September 2020. URL http://arxiv.org/abs/1802.03426. arXiv:1802.03426 [stat].

Lorenzo Nava, Maria Cuevas, Sansar Raj Meena, Filippo Catani, and Oriol Monserrat. Artisanal and Small-Scale Mine Detection in Semi-Desertic Areas by Improved U-Net. *IEEE Geoscience and Remote Sensing Letters*, 19:1–5, 2022. ISSN 1558-0571. doi: 10.1109/LGRS.2022.3220487. URL https://ieeexplore.ieee.org/abstract/document/9941123. Conference Name: IEEE Geoscience and Remote Sensing Letters.

Clement Nyamekye, Benjamin Ghansah, Emmanuel Agyapong, and Samuel Kwofie. Mapping changes in artisanal and small-scale mining (ASM) landscape using machine and deep learning algorithms. - a proxy evaluation of the 2017 ban on ASM in Ghana. *Environmental Challenges*, 3:100053, April 2021. ISSN 2667-0100. doi: 10.1016/j.envc.2021.100053. URL https://www.sciencedirect.com/science/article/pii/S2667010021000329.

Frederick Owusu-Nimo, J. Mantey, Kwabena Biritwum Nyarko, Eugene Appiah-Effah, and A. Aubynn. Spatial distribution patterns of illegal artisanal small scale gold mining (Galamsey) operations in Ghana: A focus on the Western Region. *Heliyon*, 4(2), February 2018. doi: 10.1016/j.heliyon.2018.e00534. MAG ID: 2788591747.

Jackson Simionato, Gabriel Bertani, and Liliana Sayuri Osako. Identification of artisanal mining sites in the Amazon Rainforest using Geographic Object-Based Image Analysis (GEOBIA) and Data Mining techniques. *Remote Sensing Applications: Society and Environment*, 24:100633, November 2021. ISSN 2352-9385. doi: 10.1016/j.rsase.2021.100633. URL https://www.sciencedirect.com/science/article/pii/S2352938521001695.

U.S. Geological Survey. What are rare earth elements, and why are they important?, June 2014. URL https://www.americangeosciences.org/critical-issues/faq/what-are-rare-earth-elements-and-why-are-they-important.

Global Witness. Myanmar's poisoned mountains, 2022. URL https://www.globalwitness.org/en/campaigns/natural-resource-governance/myanmars-poisoned-mountains/?gclid=Cj0KCQiAo-yfBhD_ARIsANr56g4dzQwJIj3Iep-6YEEsVLfHz-mrReBGSGN_ODvb3_qsBY1hgbd518UaAuLHEALw_wcB.

Global Witness. Fuelling the future, poisoning the present: Myanmar's rare earth boom, 2024. URL https:///en/campaigns/natural-resource-governance/fuelling-the-future-poisoning-the-present-myanmars-rare-earth-boom/.

Peipei Xia, Li Zhang, and Fanzhang Li. Learning similarity with cosine similarity ensemble. *Information Sciences*, 307:39–52, June 2015. ISSN 0020-0255. doi: 10.1016/j.ins.2015.02.024. URL https://www.sciencedirect.com/science/article/pii/S0020025515001243.

A APPENDIX

If you choose to include an appendix, please submit it as a separate PDF file.