

FLAVARS: A Multimodal Foundational Language and Vision Alignment Model for Remote Sensing

Isaac Corley^{1*} Simone Fobi Nsutezo² Anthony Ortiz² Caleb Robinson²
Rahul Dodhia² Juan M. Lavista Ferres² Peyman Najafirad¹

¹University of Texas at San Antonio

²Microsoft AI for Good Research Lab

Abstract

Remote sensing imagery is dense with objects and contextual visual information. There is a recent trend to combine paired satellite images and text captions for pre-training performant encoders for downstream tasks. However, while contrastive image-text methods like CLIP enable vision-language alignment and zero-shot classification ability, vision-only downstream performance tends to degrade compared to image-only pretraining such as MAE. In this paper, we propose FLAVARS, a pretraining method that combines the best of both contrastive learning and masked modeling, along with geospatial alignment via contrastive location encoding. We find that FLAVARS significantly outperforms a baseline of SkyCLIP for vision-only tasks such as KNN classification and semantic segmentation, +6% mIOU on SpaceNet1, while retaining the ability to perform zero-shot classification, unlike MAE pretrained methods.

1. Introduction

Multimodal learning in remote sensing with paired images and text descriptions has enabled important applications such as text-to-image retrieval, image captioning, visual-question answering, and zero-shot classification. Using language as a method of understanding satellite and aerial imagery is a clever workaround to the potentially infinite number of unique objects on the surface of the Earth. The success of pretraining methods utilizing these datasets suggests that pretraining with multimodal data can yield relevant frozen representations and weight initializations for downstream tasks. However, while multimodal pretraining provides benefits over vision-only methods, e.g. zero-shot classification, this approach generally results in a trade-off which degrades performance on visually dense tasks such as segmentation or detection. El Banani et al. [5] studied

the visual performance of a variety of foundation models and discovered that multimodal vision-language pretraining, such as CLIP [15] and SigLIP [24], performs significantly worse for vision-only tasks such as depth estimation and geometric correspondence estimation. Singh et al. [18] similarly proposed FLAVA which improved this trade-off by combining masked-image-modeling (MIM), masked-language-modeling (MLM), and contrastive learning in a unified pretraining framework.

In the field of remote sensing, imagery is generally dense with objects and visual content which can vary regionally across the globe. This creates a need for vision-language datasets to be highly detailed when describing imagery, and for pretraining to better balance visual task performance while retaining the ability to perform zero-shot classification and image-text retrieval. To better approach multimodal pretraining for remote sensing, we propose **FLAVARS** – a multimodal foundation language and visual alignment model for remote sensing. In this work, we study the following questions:

- **Does FLAVA pretraining transfer to the remote sensing domain?** We perform large-scale pretraining of FLAVA on the SkyScript dataset and compare to a SkyScript pretrained CLIP and the original FLAVA pretrained weights.
- **How can we add geospatial-awareness to FLAVA?** We add a global contrastive image-location loss to align image, text, and geospatial location. **Empirically, we find image-location alignment to be an effective addition to pretraining to improve the downstream performance of the FLAVA encoder representations.**

1.1. Related Work

Remote Sensing Image Descriptions There has recently been an increasing trend in the development of datasets containing remote sensing images and textual descriptions.

*Corresponding author: isaac.corley@utsa.edu. Work done during internship with the Microsoft AI for Good Research Lab

Model	EuroSAT	RESISC45	PatternNet	RSI-CB-256	AIRound	CyBrCT	MLRSNet	WHU-RS19	OPTIMAL-31	AID	UCMerced	FMoW
SkyCLIP	94.91	91.02	98.70	99.25	75.80	76.77	94.22	97.51	93.01	95.75	96.43	41.34
FLAVA	93.17	90.03	98.22	99.19	75.46	78.47	94.56	98.51	88.98	95.75	94.76	39.48
FLAVARS (Ours)	97.06	92.00	99.23	99.45	78.43	80.22	96.58	97.01	91.67	95.50	94.52	42.18
+ LE	97.23	92.81	99.33	99.51	79.08	81.27	96.64	97.51	91.73	95.70	95.82	43.37

Table 1. **K-Nearest Neighbor ($k = 5$) Classification results** on 12 scene recognition datasets. K-NN classification on image embeddings is an indicator of the frozen representation ability of the vision encoders. Our FLAVARS method pretrained on SkyScript outperforms SkyCLIP on 8/12 datasets. LE represents the addition of Location Encoding objective to the base FLAVA architecture resulting in FLAVARS+LE.




Figure 2. **A sample from our SkyScript-Grounded dataset.** We improve the original captions in the SkyScript dataset using GPT-4V by prompting with a caption improvement and localization instruction along with the image of interest. The grounded captions contain bounding-box pixel coordinates encompassing the objects in the image with corresponding OSM tags. As an example, we plot the resulting boxes on the sample image above.

combining OSM tags into a short sentence. We build upon this dataset by performing the following improvements to the captions:

1. We utilize GPT-4V [1] to generate detailed and descriptive paragraph-long captions of imagery by prompting the language model to utilize both the image and original caption for the improved caption generation.
2. We further incorporate visual grounding into the captions by asking GPT-4V to detect the subjects within

the caption and provide structured bounding box coordinates.

We process each image in the SkyScript dataset with the above methods to form the SkyScript-Grounded dataset which contains images, captions, geospatial coordinates, and grounded captions via bounding box annotations. A sample of the dataset is visualized in Figure 2.

2.2. FLAVARS

We take inspiration from the FLAVA pretraining framework [18] which incorporates additional masked modeling and vision-language retrieval losses in addition to contrastive learning, with the goal of improving both vision and language downstream performance and alignment.

To improve the image encoder’s geospatial representations and global awareness, we modify the FLAVA framework for the remote sensing domain by incorporating a contrastive location-image alignment. Inspired by SatCLIP [10], we utilize a location encoder initialized with the SatCLIP weights and continually pretrained using the coordinates of the SkyScript dataset. We then constrain the embeddings of the vision encoder to be jointly aligned between image, location, and text. The FLAVARS architecture is provided in Figure 1.

3. Experiments

3.1. Pretraining

We proceed by pretraining FLAVA and our FLAVARS method on the OpenAI CLIP [15] top-30% scoring subset of the SkyScript dataset [22] to fairly compare to the SkyCLIP [22] weights trained on the same subset. We note that while we introduce the SkyScript-Grounded dataset, we leave the investigation of its performance for future work.

3.2. Benchmark Datasets

We perform a thorough analysis of the performance of SkyCLIP, FLAVA, and our proposed FLAVARS methods using 12 remote sensing scene recognition datasets and the

Model	EuroSAT	RESISC45	PatternNet	RSI-CB-256	AIRound	CyBrCT	MLRSNet	WHU-RS19	OPTIMAL-31	AID	UCMerced	FMoW
SkyCLIP	37.96	65.67	69.77	44.69	50.70	26.81	60.65	88.56	83.06	69.35	73.57	18.49
FLAVA	23.06	42.19	44.41	24.81	44.28	28.57	37.44	71.64	55.38	55.45	48.81	12.93
FLAVARS (Ours)	43.07	38.16	48.49	23.96	48.66	20.67	38.04	63.18	56.99	49.30	49.05	11.51
+ LE	41.95	37.88	48.03	23.13	48.42	20.31	37.04	62.29	53.76	49.53	46.86	11.79

Table 2. **Zero-Shot Classification results** on 12 scene recognition datasets, assessing the vision-language alignment in latent space. SkyCLIP outperforms our FLAVARS method pretrained on SkyScript on 10/12 datasets, however our method still retains the ability to perform zero-shot classification and text-to-image retrieval while obtaining a better performing visual encoder. Notably, it is expected that CLIP pretraining should fare well here as alignment is the only objective during pretraining. FLAVARS, on the other hand, has multiple loss functions seeking to improve not only alignment but text and vision encoder individual performance. LE represents the addition of Location Encoding objective.

SpaceNet1 semantic segmentation dataset [19]. Furthermore, we find that several prior works appear to use different random splits for each dataset without publishing the seed used or filenames in each set. We correct this by publishing our exact splits used for training, validation, and test sets, such that the community can reproduce our work¹.

3.3. KNN Image Classification

We perform experiments by embedding the images in all image classification datasets utilizing the frozen vision encoder of each method. Results are provided in Table 1. Our results indicate that the FLAVARS pretraining significantly improves over CLIP-based pretraining by enabling the vision encoder to produce representations which can be easily classified in Euclidean space. Furthermore, we notice that adding location encoding alignment provides an additional increase in performance compared to the base FLAVA architecture.

3.4. Zero-Shot Image Classification

We perform experiments by embedding the images and textual class labels in all image classification datasets utilizing the vision and text encoders for each method. We then classify by selecting the class label with the highest cosine similarity to the image embedding. Results are provided in Table 2. Our results indicate that CLIP-based pretraining significantly improves over FLAVARS-based pretraining by producing better vision-language alignment. However, this is expected, as this objective is the primary loss used in CLIP pretraining, whereas FLAVARS has several loss functions that seek to balance and improve all representations.

3.5. Semantic Segmentation

To further assess intermediate image layer representations, we combine each pretrained vision encoder with a

Method	Model	Backbone	Pretrain Dataset	mIoU
Supervised	U-Net	ResNet-50	ImageNet	76.5
MAE	UperNet	ViT-B16	SkyScript	76.7
MAE	UperNet	ViT-B16	ImageNet	77.8
SkyCLIP	UperNet	ViT-L14	SkyScript	72.0
FLAVARS	UperNet	ViT B16	SkyScript	77.9
+ LE	UperNet	ViT B16	SkyScript	78.1

Table 3. **SpaceNet v1 Semantic Segmentation results.** Our FLAVARS method pretrained on SkyScript outperforms SkyCLIP on dense prediction tasks. We further compare to vision only pretraining methods MAE and ImageNet pretrained weights. LE represents the addition of the Location Encoding objective.

UperNet decoder [23] to form a semantic segmentation architecture. We train on the SpaceNet1 building binary segmentation dataset. Results are provided in Table 3. We further compare to other vision-only baselines not pretrained with multiple modalities – ImageNet and MAE. Our results indicate that our multimodal FLAVARS pretraining significantly improves dense vision performance over CLIP pretraining and performs well compared to other baselines.

4. Discussion & Conclusion

In this work, we proposed FLAVARS, a new multimodal foundation model for remote sensing. Building on existing models such as CLIP and FLAVA, FLAVARS demonstrates improvement in representation ability by combining masked-image & masked-language modeling with contrastive learning for remote sensing.

In Table 1 our results indicate FLAVARS is able to outperform vision-only pretrained methods like MAE [8] while retaining vision-language and vision-location alignment for zero-shot classification and retrieval. While in Table 2, CLIP outperforms FLAVARS, this is unsurprising due to the zero-shot image-text classification task being precisely the CLIP objective. FLAVARS, however, notably outperforms FLAVA in all cases for KNN classification of image embeddings and in most cases for zero-shot classification. We

¹huggingface.co/datasets/isaacacorley/FLAVARS

further find that spatial understanding can be improved by aligning image and text representations with geographic coordinates shown by the added location encoding contrastive objective. Lastly, we provide evidence there is still a noticeable trade-off in performance of vision-language pretraining across dense visual tasks, like segmentation, and multi-modal alignment capabilities, e.g. zero-shot classification. Future research is necessary to investigate this and how to systematically improve upon this trade-off.

Acknowledgements

We thank the Microsoft AI for Good Research Lab for providing the compute and resources required to perform the experiments in this paper.

References

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023. 3
- [2] Yezhen Cong, Samar Khanna, Chenlin Meng, Patrick Liu, Erik Rozi, Yutong He, Marshall Burke, David Lobell, and Stefano Ermon. Satmae: Pre-training transformers for temporal and multi-spectral satellite imagery. *Advances in Neural Information Processing Systems*, 35:197–211, 2022. 2
- [3] Isaac Corley and Peyman Najafirad. Supervising remote sensing change detection models with 3d surface semantics. In *2022 IEEE International Conference on Image Processing (ICIP)*, pages 3753–3757. IEEE, 2022. 2
- [4] Mikolaj Czerkawski and Alistair Francis. From laion-5b to laion-eo: Filtering billions of images using anchor datasets for satellite image extraction. *arXiv preprint arXiv:2309.15535*, 2023. 2
- [5] Mohamed El Banani, Amit Raj, Kevis-Kokitsi Maninis, Abhishek Kar, Yuanzhen Li, Michael Rubinstein, Deqing Sun, Leonidas Guibas, Justin Johnson, and Varun Jampani. Probing the 3d awareness of visual foundation models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21795–21806, 2024. 1
- [6] Anthony Fuller, Koreen Millard, and James Green. Croma: Remote sensing representations with contrastive radar-optical masked autoencoders. *Advances in Neural Information Processing Systems*, 36, 2024. 2
- [7] Mordechai Haklay and Patrick Weber. Openstreetmap: User-generated street maps. *IEEE Pervasive computing*, 7(4):12–18, 2008. 2
- [8] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16000–16009, 2022. 2, 4
- [9] Umangi Jain, Alex Wilson, and Varun Gulshan. Multimodal contrastive learning for remote sensing tasks. *arXiv preprint arXiv:2209.02329*, 2022. 2
- [10] Konstantin Klemmer, Esther Rolf, Caleb Robinson, Lester Mackey, and Marc Rußwurm. Satclip: Global, general-purpose location embeddings with satellite imagery. *arXiv preprint arXiv:2311.17179*, 2023. 3
- [11] Kartik Kuckreja, Muhammad Sohail Danish, Muzammal Naseer, Abhijit Das, Salman Khan, and Fahad Shahbaz Khan. Geochat: Grounded large vision-language model for remote sensing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 27831–27840, 2024. 2
- [12] Haodong Li, Haicheng Qu, and Xiaofeng Zhang. Ddfav: Remote sensing large vision language models dataset and evaluation benchmark. *arXiv preprint arXiv:2411.02733*, 2024. 2
- [13] Fan Liu, Delong Chen, Zhangqingyun Guan, Xiaocong Zhou, Jiale Zhu, and Jun Zhou. Remoteclip: A vision language foundation model for remote sensing. *arXiv preprint arXiv:2306.11029*, 2023. 2
- [14] Nathalie Pettorelli. *The normalized difference vegetation index*. Oxford University Press, USA, 2013. 2
- [15] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 1, 3
- [16] Colorado J Reed, Ritwik Gupta, Shufan Li, Sarah Brockman, Christopher Funk, Brian Clipp, Kurt Keutzer, Salvatore Candido, Matt Uyttendaele, and Trevor Darrell. Scale-mae: A scale-aware masked autoencoder for multiscale geospatial representation learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4088–4099, 2023. 2
- [17] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35:25278–25294, 2022. 2
- [18] Amanpreet Singh, Ronghang Hu, Vedanuj Goswami, Guillaume Couairon, Wojciech Galuba, Marcus Rohrbach, and Douwe Kiela. Flava: A foundational language and vision alignment model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15638–15650, 2022. 1, 3
- [19] Adam Van Etten, Dave Lindenbaum, and Todd M Bacastow. Spacenet: A remote sensing dataset and challenge series. *arXiv preprint arXiv:1807.01232*, 2018. 4
- [20] Yi Wang, Nassim Ait Ali Braham, Zhitong Xiong, Chenying Liu, Conrad M Albrecht, and Xiao Xiang Zhu. Ssl4eo-s12: A large-scale multi-modal, multi-temporal dataset for self-supervised learning in earth observation. *arXiv preprint arXiv:2211.07044*, 2022. 2
- [21] Yi Wang, Hugo Hernández Hernández, Conrad M Albrecht, and Xiao Xiang Zhu. Feature guided masked autoencoder for self-supervised learning in remote sensing. *arXiv preprint arXiv:2310.18653*, 2023. 2

- [22] Zhecheng Wang, Rajanie Prabha, Tianyuan Huang, Jiajun Wu, and Ram Rajagopal. Skyscript: A large and semantically diverse vision-language dataset for remote sensing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 5805–5813, 2024. [2](#), [3](#)
- [23] Tete Xiao, Yingcheng Liu, Bolei Zhou, Yuning Jiang, and Jian Sun. Unified perceptual parsing for scene understanding. In *Proceedings of the European conference on computer vision (ECCV)*, pages 418–434, 2018. [4](#)
- [24] Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11975–11986, 2023. [1](#)
- [25] Yang Zhan, Zhitong Xiong, and Yuan Yuan. Rsvg: Exploring data and models for visual grounding on remote sensing data. *IEEE Transactions on Geoscience and Remote Sensing*, 61:1–13, 2023. [2](#)
- [26] Zilun Zhang, Tiancheng Zhao, Yulong Guo, and Jianwei Yin. Rs5m: A large scale vision-language dataset for remote sensing vision-language foundation model. *arXiv preprint arXiv:2306.11300*, 2023. [2](#)