Distilling geospatial foundation models in Earth observation

Internship for Master 2 or École d'Ingénieur student
(With an opportunity to pursue a PhD)

Expected starting: February/March 2026 (6 months)

Keywords

Earth observation, remote sensing, foundation models, knowledge distillation

Context and Objectives

Context. In recent years, deep learning (DL) has found success in various machine learning and computer vision domains, including remote sensing (RS). Despite this, applying DL to real-world scenarios using VHR satellite images for operational purposes faces numerous challenges. The foremost challenge involves the difficulty of annotating domain-specific data, particularly in EO applications that demand expert knowledge. Generating precise and comprehensive labeled datasets for training deep models is a time-consuming and expensive endeavor. For instance, in rapid disaster mapping, acquiring accurate labels is nearly unfeasible due to the infrequent and unique nature of catastrophic events. Another challenge comes from the significant domain shifts inherent in RS data, arising from diverse sensor characteristics (i.e., spatial resolution and spectral bands) and varied acquisition conditions. Lastly, the exploding size of deep models, with millions (or even billions) of parameters, raises concerns. Not only do they demand substantial computational and storage resources, but they also cause negative environmental issues.

Therefore, designing efficient models while maintaining high accuracy becomes essential in every EO application to reduce energy cost and more importantly, to minimize the environmental impact. These models are expected to be reusable or transferred with low resources. Recent studies have showed that by leveraging self-supervised learning (SSL) on the abundance of multi-source unlabeled data, geospatial foundation models (FMs) [1] have started their era by providing high performance on a wide range of downstream tasks. However, current trends of FMs in EO only focus on creating large vision FMs using substantial multi-source images such as SkySense [2], AnySat [3], etc. (see a survey in [1]). These models demand significant computational resources for training and deployment. Therefore, developing resource-efficient models in the context of EO is imperative to mitigate environmental concerns in the future [4]. One way to address the efficiency concern of FMs is the usage of model compression algorithms, such as *knowledge distillation* [5]. This technique is widely used in deep learning to reduce the size of a model while keeping its performance high. In the field of EO, knowledge distillation of FMs has been used to assist the training of a specialist model [6], or to obtain smaller size FMs [7], but its potential remain largely unexplored [8].

Objectives of this work. This internship topic aims to distill compact and lightweight models from large geospatial FMs while maintaining their high performance on various downstream EO tasks. The idea is to perform distillation not only from a single FM, but on several FMs in order to achieve their complementary capacity to better deal with various EO tasks. The benchmarking and evaluation will follow the well-established protocols in GeoBench [9] and Pangaea [10]. The proposed topic is expected to provide cutting-edge solutions for large-scale EO applications using multi-source satellite images in the context of sustainable development.

Work Planning

To address the aforementioned objectives, a tentative plan is given below.

- Literature review of geospatial/remote sensing FMs and knowledge distillation techniques.
- Identification of datasets and tasks related to disaster mapping using satellite images.
- Implementation of knowledge distillation from single- and multi-teacher FMs, to small specialist models and small FMs.
- Evaluation and comparison of the proposed approaches.

Candidate Profile

We are looking for a candidate

- who is a Master 2, École d'Ingénieur or equivalent student in computer science or applied mathematics with excellent academic track;
- has a strong background in computer science, machine learning, and/or computer vision and image processing;
- has excellent programming in Python and is familiar with PyTorch;
- potentially has experience with remote sensing, or at least a keen interest in the application;
- has excellent communication skills in French or English;
- has a keen interest in research — please note that **there is a funded opportunity to pursue a PhD related to this topic**.

Organisation

The master thesis will start around February / March and last for a period of 5 to 6 months. The research will be done at Université Bretagne Sud in Vannes. The expected intern will join the OBELIX research group[1] from IRISA (UMR 6074) located at the Université Bretagne Sud in Vannes, France. He/She will be supervised by:

- Mr. Pierre Adorni, PhD student UBS/IRISA,
- Dr. Minh-Tan Pham[2], Assoc. Professor (Maître de Conférences, HDR) UBS/IRISA

This internship is part of the DREAMS project (ANR JCJC) funded by the French research agency (ANR).

[1] www.irisa.fr/obelix
[2] https://sites.google.com/site/mtanpham89

Application

Send your **CV + Motivation letter + transcripts for the two last years** to

- `pierre.adorni@irisa.fr`
- `minh-tan.pham@irisa.fr`

Deadline: **15 December 2025**. Selected candidates will be contacted for interview.

References

[1] S. Lu, J. Guo, J. R. Zimmer-Dauphinee, J. M. Nieusma, X. Wang, S. A. Wernke, Y. Huo, *et al.*, "Vision foundation models in remote sensing: A survey," *IEEE Geoscience and Remote Sensing Magazine*, 2025.

[2] X. Guo, J. Lao, B. Dang, Y. Zhang, L. Yu, L. Ru, L. Zhong, Z. Huang, K. Wu, D. Hu, *et al.*, "Skysense: A multi-modal remote sensing foundation model towards universal interpretation for earth observation imagery," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 27672–27683, 2024.

[3] G. Astruc, N. Gonthier, C. Mallet, and L. Landrieu, "Anysat: One earth observation model for many resolutions, scales, and modalities," in *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 19530–19540, 2025.

[4] P. Adorni, M.-T. Pham, S. May, and S. Lefèvre, "Towards efficient benchmarking of foundation models in remote sensing: A capabilities encoding approach," in *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 3096–3106, 2025.

[5] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *arXiv preprint arXiv:1503.02531*, 2015.

[6] T. Gao, W. Ao, X.-a. Wang, Y. Zhao, P. Ma, M. Xie, H. Fu, J. Ren, and Z. Gao, "Enrich distill and fuse: Generalized few-shot semantic segmentation in remote sensing leveraging foundation model's assistance," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2771–2780, 2024.

[7] O. Siméoni, H. V. Vo, M. Seitzer, F. Baldassarre, M. Oquab, C. Jose, V. Khalidov, M. Szafraniec, S. Yi, M. Ramamonjisoa, *et al.*, "Dinov3," *arXiv preprint arXiv:2508.10104*, 2025.

[8] Y. Himeur, N. Aburaed, O. Elharrouss, I. Varlamis, S. Atalla, W. Mansoor, and H. Al-Ahmad, "Applications of knowledge distillation in remote sensing: A survey," *Information Fusion*, vol. 115, p. 102742, 2025.

[9] A. Lacoste, N. Lehmann, P. Rodriguez, E. Sherwin, H. Kerner, B. Lütjens, J. Irvin, D. Dao, H. Alemohammad, A. Drouin, *et al.*, "Geo-bench: Toward foundation models for earth monitoring," *Advances in Neural Information Processing Systems*, vol. 36, pp. 51080–51093, 2023.

[10] V. Marsocci, Y. Jia, G. L. Bellier, D. Kerekes, L. Zeng, S. Hafner, S. Gerard, E. Brune, R. Yadav, A. Shibli, *et al.*, "Pangaea: A global and inclusive benchmark for geospatial foundation models," *arXiv preprint arXiv:2412.04204*, 2024.