Geospatial Artificial Intelligence (GeoAI) Insights
ISSN: XXXX Vol. 01, Issue: 01, 16-29, 2026
DOI: XXXX

Few-Shot Learning for Land Cover Classification Using Sentinel-2 Imagery: A Lightweight GeoAI Framework

Arpit Sharma[1,*], Imane AITOUHANNI[2], Nouha Belmahi[3], Yassine Mouniane[4], Mohamed El Bakkali[5]

[1]National Institute of Pharmaceutical Education and Research, Mohali, India
[2]Mohammed V University, ENSIAS, SSLAB, Rabat, Morocco
[3]Faculty of Economics and Management, Ibn Tofail University, Morocco
[4]Natural Resources and Sustainable Development Laboratory, Faculty of Sciences, Ibn Tofail University, B.P. 242, Kenitra, Morocco
[5]Higher Institute of Nursing Professions and Health Technologies, Rabat, Kenitra Annex, Morocco

** Corresponding author: sarpit909@gmail.com*

Abstract: Recent advances in satellite remote sensing have led to the widespread availability of high-resolution Earth observation data. However, the performance of deep learning models in land cover classification often depends on large volumes of annotated data, which are costly and time-consuming to obtain. This study applies a lightweight GeoAI approach for land cover classification using Sentinel-2 imagery under limited training data conditions. We implement a lightweight GeoAI framework utilizing prototypical networks with pre-trained convolutional backbones to classify land cover types under low-data regimes. Experiments on a subset of the EuroSAT dataset demonstrate that our approach achieves competitive accuracy, even with as few as one to five labeled examples per class. This work highlights the potential of FSL techniques in reducing labeling costs and enabling rapid deployment of land cover models in data-scarce environments.

Keywords: Few-shot learning, Land cover classification, Remote sensing, Sentinel-2 imagery, EuroSAT dataset, Prototypical networks, GeoAI, Meta-learning.

1 Introduction

Identification and mapping of land cover types is a key component of remote sensing, which is used for applications in agriculture, urban planning, forestry, environmental monitoring, and many others. Conventional supervised learning schemes for analyzing satellite images require large, labeled datasets, which can often be expensive or impossible to collect, for example, for a first-time observed area or an under-represented geography [1]. The advent of satellite missions, such as ESA's Sentinel-2, that acquire high-resolution multispectral imagery and have high revisit frequencies, means the possibilities for real-time, large-scale

Received: 30 Aug. 2025; Revised: 05 Nov. 2025; Accepted: 18 Dec. 2025; Published: 25 Oct. 2026

© 2026 The Author(s). Published by Science Transactions. License: CC BY 4.0

land cover mapping have never been wider [2]. However, the annotation bottleneck remains a critical challenge for scalable geospatial intelligence systems. Few-Shot Learning (FSL) refers to a class of machine learning techniques designed to recognize new categories using only a handful of labeled samples per class. FSL uses previous knowledge and meta-learning techniques to quickly adapt models to new tasks,especially useful in geospatial tasks where labeled data are scarce. FSL provides a potential solution to this problem because it can generalize to unfamiliar classes from only a few labeled examples.FSL was first examined in the fields of computer vision and natural language processing, but in the last few years, it has been modified to conduct a revised form for remote-sensing tasks. FSL methods usually use meta-learning techniques to prepare models for fast adaptation to new tasks [3].This study aims to evaluate the potential use of Few-Shot Learning (FSL) for Sentinel-2 data and develop a lightweight GeoAI framework for land cover mapping tasks. The main contributions are: (i) the development of a prototypical network-based model for geospatial data, (ii) its experimental assessment in a few-shot or limited-data scenario using the EuroSAT dataset, and (iii) discussions about computational efficiency and the possibility of real-time use.The main contributions to this study are its original contributions concerning the development of a lightweight Few-Shot Learning (FSL) framework for the geospatial task of land cover classification, its empirical assessment in a few-shot or limited annotation scenario using real Sentinel-2 data, and the exploration of the framework's potential use in a constrained resource or a rapidly changing geospatial context.The rest of the paper is structured as follows. Section 2 reviews the related work on deep learning and FSL in remote sensing. Section 3 states the proposed methodology. Section 4 describes the experimental setup and results. Section 5 provides a discussion of findings, and Section 6 concludes with future directions.

2 Related Work

The remote sensing community has traditionally prioritized land cover classification from satellite images. Modern deep learning methods, and more specifically convolutional neural networks (CNNs), have achieved high performance on large annotated datasets for tasks such as scene classification and semantic segmentation. Publicly available datasets like EuroSAT [4], BigEarthNet [5], and NWPU-RESISC45 [6] have led to significant advancements in this area due to standardized benchmarks and geographic coverage. However, each of these datasets requires considerable human annotation, limiting scalability and transferability to new regions or classes. Prototypical Networks are among the effective methods in this space, computing prototypes for each class in a learned embedding space and classifying based on distance [7]. Other architectures, such as matching networks and relation networks,have also been demonstrated to be effective on low-data vision tasks [8]. Recently, few-shot learning techniques have expanded into remote-sensing-related contexts. For example, Liu et al. [9] investigated few-shot learning approaches for remote sensing image interpretation, while Broni-Bediako et al. [10] explored a generalized few-shot semantic segmentation approach. Zhan et al. developed a deep Siamese CNN framework for change detection, which is another geospatial task that, like land cover classification, relies on a limited number of labeled samples[11]. Nevertheless, the majority of recent studies in remote sensing contexts using few-shot learning approaches are limited to hyperspectral data or controlled experiments.

The potential for few-shot models to apply to multispectral datasets, widely available from sources such as Sentinel-2 imagery, has remained comparatively less explored. Additionally, most work in this area has not fully addressed the trade-offs between model complexity, computational cost, and feasibility in real-world environments with limited resources. In this paper, we address this issue by showing that prototypical networks in conjunction with pre-trained CNNs can be used to create a lightweight and simple few-shot learning framework to perform land cover classification using Sentinel-2 imagery.

3 Methodology

This section outlines the overall framework of our approach, including the dataset used, the few-shot learning setup, the architecture of the prototypical network, and the implementation details. Figure 1 shows the proposed workflow to clarify the various steps taken in this study. Starting with Sentinel-2 RGB images as input, features from the images were extracted using a pre-trained ResNet-18 backbone. The extracted image embeddings were used in a Few-Shot Learning (FSL) model based on prototypical networks to build class prototypes from support samples.Next, query samples were classified by finding the nearest prototype in the embedding space, producing the final land cover map output.

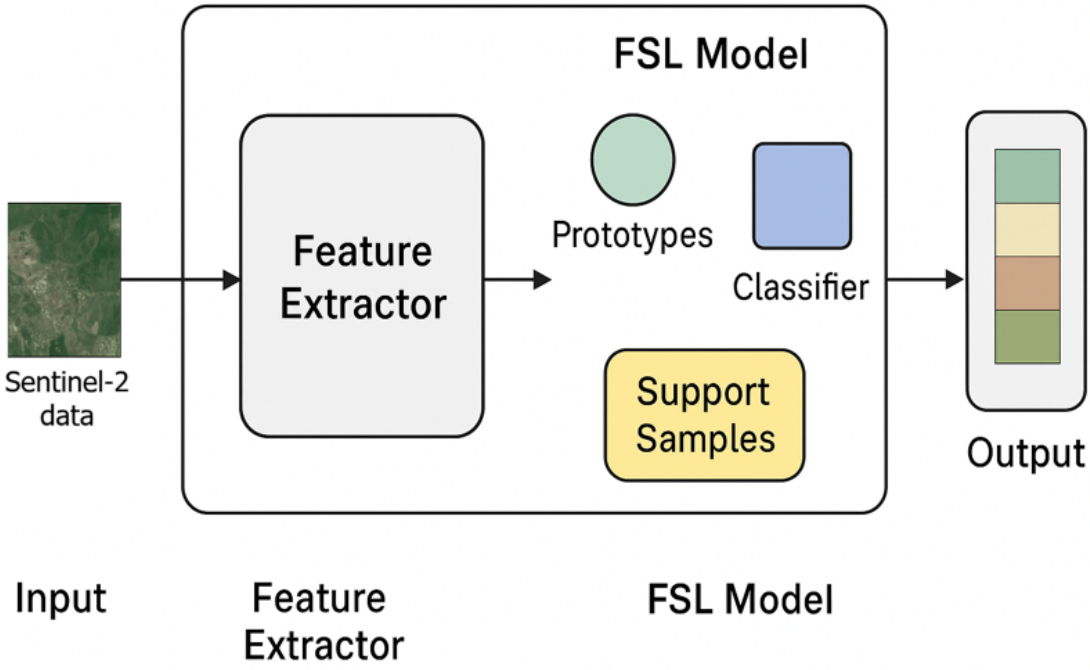


Fig. 1. Workflow of the proposed Few-Shot Learning framework for land cover classification using Sentinel-2 imagery. The model integrates a pre-trained feature extractor with prototype-based classification to enable accurate predictions under limited data conditions.

3.1 Dataset and Preprocessing

The study area corresponds to the EuroSAT dataset [4], which consists of Sentinel-2 image patches collected from multiple regions across Europe, approximately between latitudes 35° N and 60° N, and longitudes −10° E and 30° E, covering 34 countriesFigure 2.

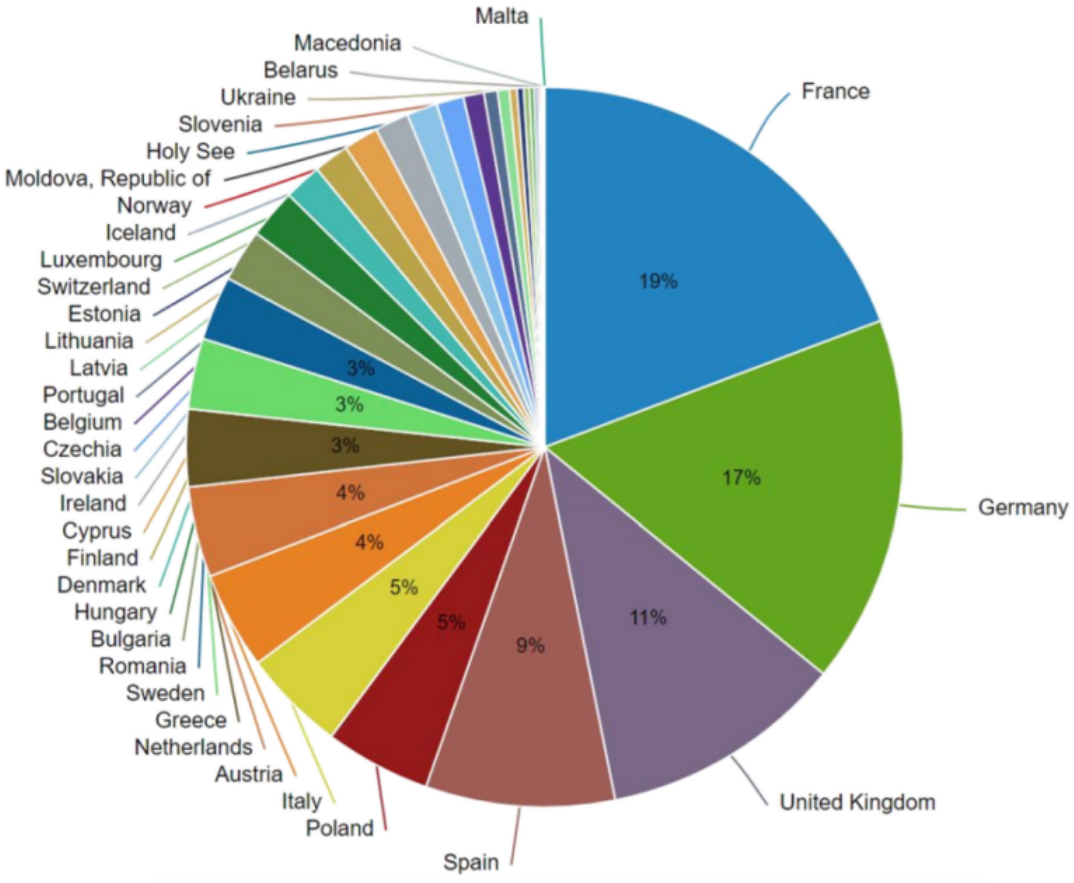


Fig. 2. EuroSAT dataset distribution. The geo-referenced images are distributed across Europe. The distribution is influenced by the number of represented cities per country in the European Urban Atlas.

The dataset comprises 27,000 geo-referenced and labeled images in both RGB and multi-spectral formats, representing ten land use and land cover classes such as Residential, Forest, Annual Crop, and Industrial areas. Figure 3 presents examples of land use and land cover classes.

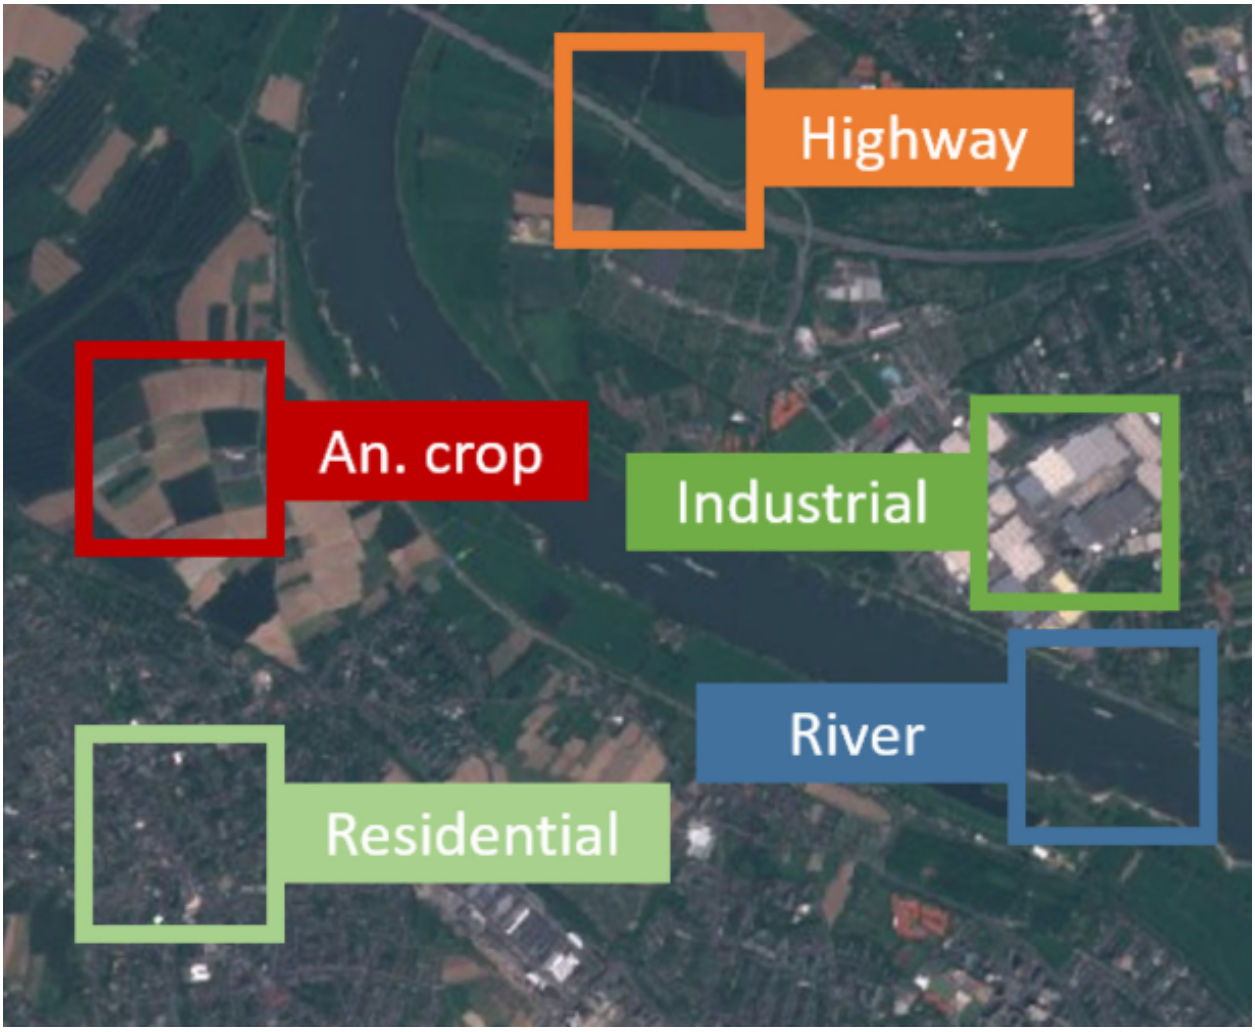


Fig. 3. Land use and land cover classification based on Sentinel-2 satellite images. Patches are extracted for class identification. This visualization highlights the classes: Annual Crop, River, Highway, Industrial Buildings, and Residential Buildings.

We use the RGB subset for simplicity and so that the results can be reproduced. Each image is 64×64 pixels and was normalized using the ImageNet mean and standard deviation for consistency with the pre-trained convolutional backbones. To simulate a few-shot learning setting, we randomly selected N classes with K support samples for each class, and some number of query samples to evaluate the model.

3.2 Few-Shot Learning Setup

We adopt an N-way K-shot episodic training framework, where the model learns to classify among N classes using K labeled examples per class. During training, each episode simulates a classification task consisting of:

- A support set $\mathbf{S} = \{(x_i, y_i)\}_{i=1}^{NK}$ with K examples from each of the N classes.
- A query set $\mathbf{Q}$ containing unlabeled examples used for evaluation.

For example, a 5-way 1-shot episode consists of 5 labeled samples (one per class) in the support set and 15 query samples (3 per class).

3.3 Prototypical Networks

The core of our framework is the Prototypical Network [7], which maps input images to an embedding space where classification is performed based on distances to class prototypes. Let $f_\theta(x)$ denote the embedding function, such as a ResNet encoder. For each class c, we compute a prototype p_c as the mean of the embedded support examples in that class:

$$p_c = \frac{1}{K} \sum_{(x_i, y_i) \in S_c} f_\theta(x_i) \tag{1}$$

A query example x is classified based on the Euclidean distance to each prototype:

$$\hat{y} = \arg\min_c \|f_\theta(x) - p_c\|^2 \tag{2}$$

This approach is computationally efficient, interpretable, and effective for few-shot classification.

3.4 Feature Extractor

In the experiment, we use ResNet-18 initialized with weights trained on ImageNet in order to perform feature extraction.ConvNets have good generalization properties with respect to large datasets and are especially useful in situations involving limited data in the target domain. ResNet-18 is an architecture that consists of 18 convolutional layers and includes residual connections to address the gradient vanishing problem, thus allowing for deeper networks without difficulty. We remove the final fully connected classification layer, retaining all the convolutional blocks, including the global average pooling layer.Each input image is resized to 64×64 pixels, thereby mapping to a 512-dimensional feature vector. To keep the model lightweight for few-shot situations and reduce overfitting, we freeze the weights of the pre-trained ResNet-18 while training. This approach speeds up the training process while benefiting from general feature representations that are transferred from ImageNet, even though the domain may differ. If additional computational resources and domain-specific data are available, the embedding network could optionally be fine-tuned for improved performance. The features extracted from the support set are then used to compute class prototypes, while the features extracted from the query set are used for classification using the distance metrics outlined in Section 3.3.

3.5 Training Protocol

The prototypical network is trained using the episodic training paradigm, which is similar to the few-shot testing scenario.Rather than the standard mini-batch training process, we create a sequence of classification tasks, or episodes, that utilize a support set and a query set of samples. The episodic training format allows the model to learn a generalizable inductive bias, which then enables the model to quickly adapt to new tasks with less training data.

For each training episode, we randomly sample N classes from the training set and sample K support examples from these N classes. We then use additional examples from the same N classes to form the query set.For training, we used 5-way classification tasks with either a 1-shot (1 support image per class) or a 5-shot (5 support images per class) configuration of samples in each episode to evaluate each model under different training data availability situations.We train the model using the Adam optimizer with a learning rate of 1×10^{-3} and a batch size equivalent to one episode.The loss function is the negative log-likelihood derived from the softmax over distances between query embeddings and class prototypes. Training is carried out for 1,000 episodes, with early stopping based on validation accuracy to prevent overfitting.To promote consistent results and solid evaluation, we use a fixed random seed for episode generation.The training set is distinct from the test classes to ensure a cross-class generalization requirement, which is a defining attribute of few-shot learning.

3.6 Evaluation Metrics

The evaluation of the model's performance is based on mean classification accuracy from 600 test episodes randomly sampled from the complete list of episodes.Each episode had unseen classes during training, thus providing a few-shot testing environment.For each test episode query image, the model predicts a class based on the distance between the query image and all class prototypes. Furthermore, we report the standard deviation across test episodes in addition to mean accuracy to account for the model's stability and robustness across random configurations of classes. This form of evaluation is more informative than accuracy from a single sampled episode, as it illustrates how stable and resilient the model is across different episodes representing random sampling.Finally, we also computed confusion matrices to assess the model's performance on each class. The confusion matrices were useful for visualizing systematic misclassification, especially where classes contain closely similar visual characteristics, such as Forest versus Herbaceous Vegetation.

4 Experiments and Results

4.1 Accuracy vs. Number of Shots

In order to study the effect of having labeled data on model performance, we evaluated our experiment using different K-shot settings.The results are shown in Figure 4. We can see that the accuracy increases consistently as the number of support examples per class increases.With the 1-shot setting, the model achieves an accuracy of 58.2%, indicating the difficulty of training with a single example per class. However, the performance improves to 76.8% with 5 shots and reaches 83.5% with the 10-shot setting.This trend is consistent with the expected behavior of few-shot learning.When more labeled data points are available for each class, the prototypes can be estimated more accurately, resulting in more accurate distance-based classification.The standard deviation, which is also plotted in Figure 4 as error bars, decreases as the number of shots increases, suggesting that the predictions become more consistent and stable.

4.2 Class-wise Analysis via Confusion Matrix

We analyzed the confusion matrix for the 5-way 5-shot classification task shown in Figure 5 to better understand which classes were confused more often.The confusion matrix showed that Annual Crop, Forest, and Residential classes performed well, showing high true positive counts.Confusion between Industrial and Annual Crop classes is evident, possibly due to both classes having similar spatial textures and colors in some of the Sentinel-2 images.Forest and Residential classes were sometimes confused as well, possibly due to urban-forest boundaries or mixed land use in the imagery.

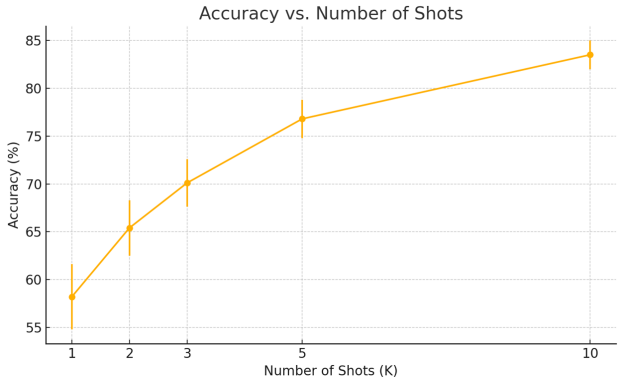


Fig. 4. Accuracy versus number of shots (K) in 5-way classification.

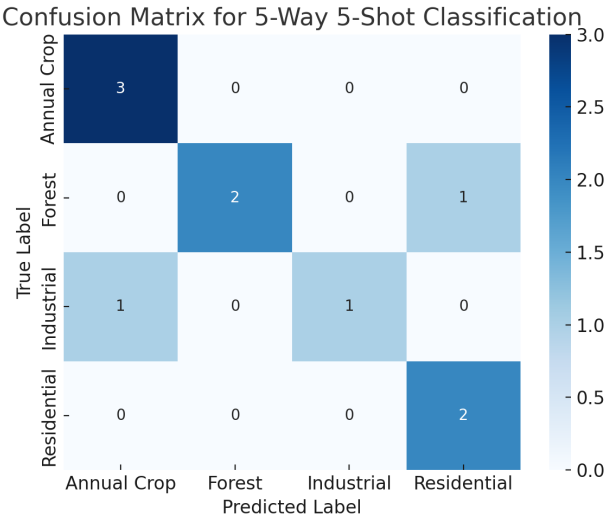


Fig. 5. Confusion matrix for 5-way 5-shot classification showing prediction distribution across four land cover types.

4.3 Spatial Interpretation

Figure 6 provides an enhanced spatial interpretation of the model's predictions through a map of land cover classification using Sentinel-2 RGB imagery. The map depicts the spatial arrangement of predicted land cover classes, specifically Residential, Forest, Annual Crop, and Industrial zones over matched areas of the EuroSAT dataset. Overall, the results demonstrate that the proposed Few-Shot Learning (FSL) framework captures cohesive spatial patterns reflective of the real land cover boundaries in the study area. The model is able to differentiate forested areas from urban and industrial areas, while agricultural lands are primarily positioned in rural areas at lower elevations. Furthermore, mixed class transitions from urban-rural boundaries are visible on the classified map, which is consistent with trends presented in the confusion matrix. In summary, the spatial representation demonstrates the potential of the proposed GeoAI model to support large-scale, data-efficient land cover mapping using a limited number of annotated samples.

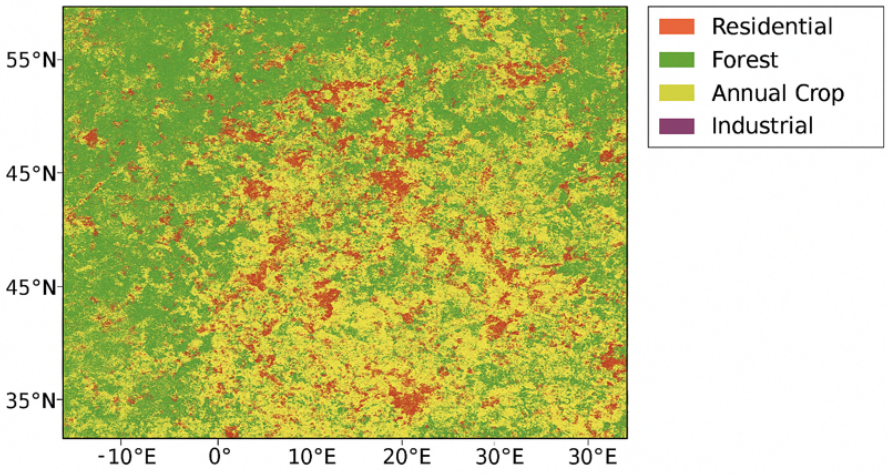


Fig. 6. Predicted land cover classification map using Sentinel-2 RGB imagery.

The map displays the distribution of land-cover types, including Residential, Forest, Annual Crop, and Industrial categories, across the area included in the EuroSAT study, approximately 35° N to 60° N and −10° E to 30° E.The classification demonstrates the potential of the Few-Shot Learning framework to characterize different types of land cover while preserving an approximation of spatial coherence across the study area. The EuroSAT data display considerable spectral and spatial diversity within the classes.Built-up classes, including Residential and Industrial, are notably different from natural classes such as Forest and Herbaceous Vegetation in terms of textural contrast. Agricultural classes, including Annual and Permanent Crops, tend to show different spectral patterns based on phenological stages, which add intra-class variability and make few-shot learning especially complex.

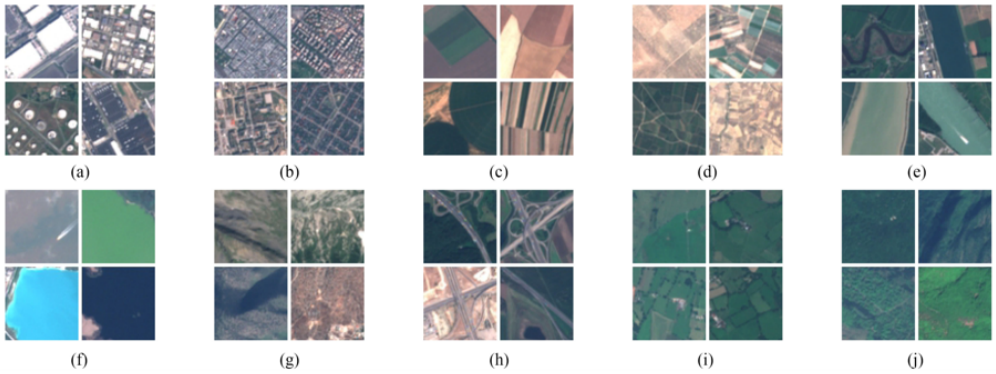


Fig. 7. Representative examples of the ten land cover classes included in the EuroSAT dataset. Each patch measures 64 × 64 pixels and corresponds to a distinct Sentinel-2 land cover category. (a) Industrial buildings. (b) Residential buildings. (c) Annual crop. (d) Permanent crop. (e) River. (f) Sea and lake. (g) Herbaceous vegetation. (h) Highway. (i) Pasture. (j) Forest. Adapted from Helber et al. [4].

5 Discussion

The experimental results presented in Section 4 illustrate the performance of Few-Shot Learning (FSL), specifically prototypical networks, for land cover classification using Sentinel-2 data. The results in Figure 4 show that model accuracy increases with the addition of labeled support examples, demonstrating the effect of additional samples on improving the estimation of prototypes and reducing intra-class variance in the embedding space. Despite training with limited data, the model was able to reach an accuracy of 58.2% in the 1-shot setting, indicating that pre-trained convolutional features can retain discriminative power with very limited data.This is an important capability in real-world GeoAI applications where models need to adapt quickly with limited annotation overhead. The visible spectral ranges in these EuroSAT classes reflect the inherent spectral and spatial variability of the original Sentinel-2 image data. Built-up classes, such as Residential Buildings and Industrial Buildings, are more complex and have sharper edges, whereas natural classes such as Forest and Herbaceous Vegetation tend to be more homogeneous and spectrally stable. Agricultural classes, including Annual Crop and Permanent Crop, have higher intra-class variability due to crop phenology and seasonality, compounding the difficulty of few-shot generalization. The spatial extent of the dataset, shown in Figure 2, spans 34 European countries between latitudes 35° N and 60° N and longitudes −10° E and 30° E, adding variability to the examples used for model development. The confusion matrix shown in Figure 5 provides additional details regarding class-specific performance.The network achieves high recognition performance for visually distinct classes such as Forest, Sea & Lake, and Residential.However, there is moderate confusion between the Industrial and Annual Crop classes. The confusion between Industrial and Annual Crop classes may stem from a combination of color distribution and mixed land uses at urban-rural interfaces. Helber et al. [4] also reported challenges involving visually similar classes, suggesting that spectral ambiguities at the patch level may persist even in high-resolution imagery. These findings also highlight one limitation of prototype-based approaches: small changes in spectral characteristics may lead to misclas-

sification when class centroids are relatively close in theembedding space.Spatial-attention mechanisms, spectral-specific features, and adaptive measures could contribute to improving prototype-based classification.In addition to RGB analysis, Figure 8, modified from Helber et al. [4], shows that additional spectral bands on Sentinel-2, particularly Red Edge 1 (B05) and SWIR2 (B12), can provide discriminative information for distinguishing vegetation and soil.Our current framework was built using only the RGB subset for computational simplicity, but future iterations including the additional bands may increase classification separability, particularly between agricultural and built-up land uses.This modification could provide richer representations of biophysical attributes such as canopy density, chlorophyll content, and surface moisture.

The amalgamation of visual, spatial, and spectral evidence supports the idea that lightweight few-shot frameworks can overcome limitations associated with both data scarcity and large-scale geospatial analyses. The efficiency and flexibility of the model can be applied to situations such as rapid post-disaster mapping, land-use change monitoring, and updates to regional inventories that rely on limited labeled data.

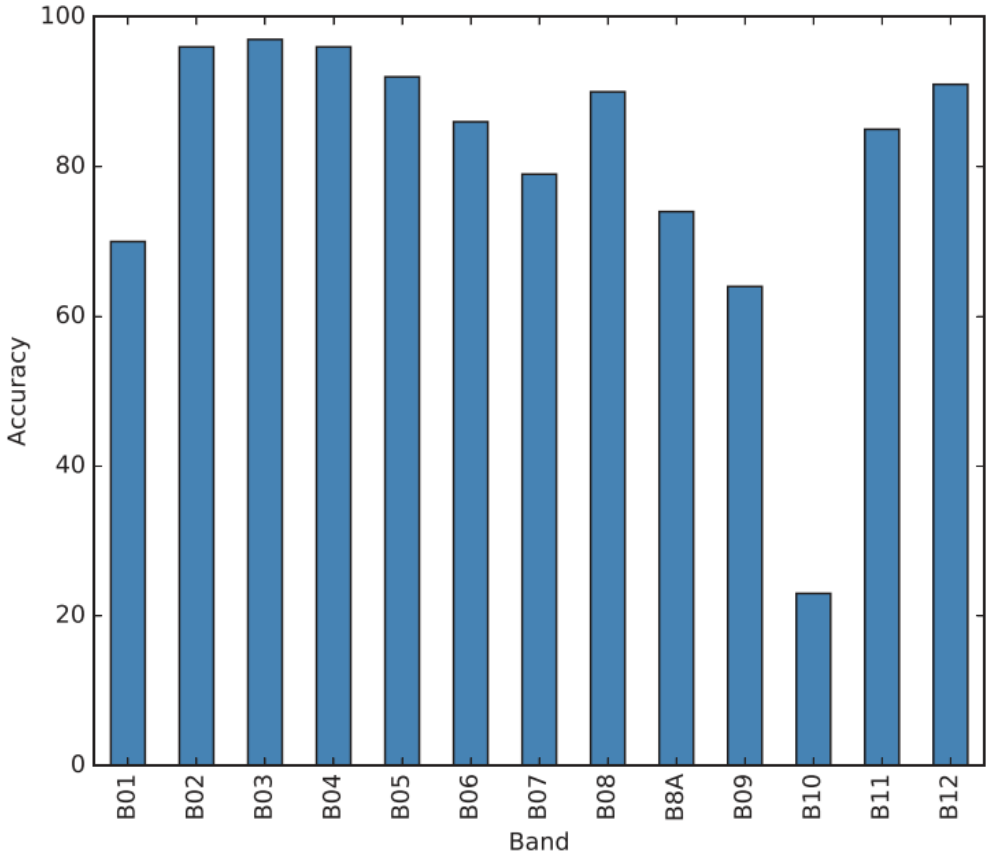


Fig. 8. Comparison of classification accuracy across Sentinel-2 spectral bands, adapted from Helber et al. [4]. Red Edge 1 (B05) and SWIR2 (B12) bands demonstrate high discriminative power beyond the RGB configuration used in this study.

Future research directions may include hybrid FSL-multispectral models, embedding strategies through domain adaptation, or temporal generalization across multiple acquisi-

tion dates of Sentinel-2 images. Future developments will advance Few-Shot Learning as a relevan component of GeoAI frameworks for sustainable land management. The framework evaluated in this study functions effectively under few-shot conditions but has several limitations that should be addressed in future studies. First, while the RGB subset of EuroSAT is computationally efficient, it limits the spectral richness available from the full Sentinel-2 dataset. Multispectral or temporal features could provide additional information on vegetation indices, soil moisture, and built materials, potentially improving class separability in complicated or transitional landscapes. Second, the present application is based on a frozen ResNet-18 backbone pre-trained on ImageNet. While this approach increases stability and helps mitigate the risk of overfitting, it can also restrict the classifier's capacity to learn domain-specific spectral and spatial features unique to remote sensing images. Adapting the feature extractor through fine-tuning or designing a dual-branch architecture consisting of RGB and spectral bands may improve adaptation to geospatial settings. Third, assessment was limited to the EuroSAT dataset which, although well distributed spatially, mainly sampled landscapes within Europe. Incorporating additional datasets such as BigEarthNet, DeepGlobe, or SEN12MS could provide an assessment of model transferability across biomes, climatic zones, and acquisition conditions. The use of Euclidean distance as the similarity metric in the prototypical network may also be a simplified assumption for more complex relationships between classes. Further work should explore advanced distance functions, prototype refinement using attention, or transformer-based meta-learning to better discriminate features for land cover categories in extreme low-shot situations. Introducing uncertainty estimation could also facilitate decision support for practical GeoAI systems by informing users about model confidence across heterogeneous spatial domains.Overall, addressing these limitations can contribute to more robust, interpretable, and scalable few-shot learning solutions for satellite-based land cover classification and mapping.

6 Conclusions

This study developed a lightweight few-shot learning framework for land cover classification using Sentinel-2 data and evaluated the model's performance on the EuroSAT dataset under various conditions of low data availability. More specifically, we used prototypical networks with a fixed ResNet-18 structure. We shown that useful performance in few-shot contexts can be achieved with limited modeling overhead and that such frameworks can be applied to GeoAI workflows, especially in scenarios where labeled data are scarce or rapid implementation is required. The model demonstrated potential for generalization to unseen categories, even with only one labeled sample per category, which supports the development of scalable and data-efficient land cover mapping approaches.

Looking ahead, we plan to extend this work in several directions:

- Incorporate multispectral inputs to better leverage Sentinel-2's rich band structure.
- Explore meta-learning-based fine-tuning strategies to further adapt the model to domain-specific variations.
- Evaluate the framework on cross-region generalization tasks to assess robustness in different geographical settings.

Ultimately, this study contributes to the growing body of research in GeoAI by demonstrating the applicability of few-shot learning to geospatial classification challenges involving limited labeled data.

Author Contributions (Compulsory)

Arpit Sharma conceived the study, implemented the Few-Shot Learning model, and wrote the initial draft of the manuscript. **Imane Aitouhanni** contributed to the experimental design, data preprocessing, and validation of the methodology. **Nouha Belmahi** assisted with visualization, statistical analysis, and interpretation of the results. **Yassine Mouniane** supervised the overall research, provided resources, and revised the manuscript for intellectual content and clarity. **El Bakkali Mohamed** contributed to the methodological refinement,critical review, and final approval of the submitted version.

Conflict of Interest

The authors declare no conflict of interest.

Research Involving Human and/or Animals

Not Applicable.

Data Availability

Not applicable.

Funding

The authors received no financial support for this research.

AI Use Disclosure

The author(s) confirm that generative AI tools were not used in the preparation of this article.

References

[1] Yue, J., Fang, L., Ghamisi, P., Xie, W., Li, J., Chanussot, J., Plaza, A. (2022). Optical remote sensing image understanding with weak supervision: Concepts, methods, and perspectives. IEEE Geoscience and Remote Sensing Magazine, 10(2), 250-269. https://doi.org/10.1109/MGRS.2022.3161377

[2] Mazzia, V., Khaliq, A., Chiaberge, M. (2019). Improvement in land cover and crop classification based on temporal features learning from Sentinel-2 data using recurrent-convolutional neural network (R-CNN). Applied Sciences, 10(1), 238.

https://doi.org/10.3390/app10010238

[3] Rußwurm, M., Wang, S., Körner, M., Lobell, D. (2020, June). Meta-learning for few-shot land cover classification. In 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW) (pp. 788-796). IEEE. https://arxiv.org/abs/2004.13390

[4] Helber, P., Bischke, B., Dengel, A., Borth, D. (2019). Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, 12(7), 2217-2226. https://doi.org/10.1109/JSTARS.2019.2918242

[5] Sumbul, G., Charfuelan, M., Demir, B., Markl, V. (2019, July). Bigearthnet: A large-scale benchmark archive for remote sensing image understanding. In IGARSS 2019-2019 IEEE international geoscience and remote sensing symposium (pp. 5901-5904). IEEE. https://doi.org/10.1109/IGARSS.2019.8900532

[6] Cheng, G., Han, J., Lu, X. (2017). Remote sensing image scene classification: Benchmark and state of the art. Proceedings of the IEEE, 105(10), 1865-1883. https://doi.org/10.1109/JPROC.2017.2675998

[7] Omololu, A. A., Olaniyi, A. J., & Olayinka, O. O. (2025). Application of few-shot learning using prototypical networks for classification of neglected tropical diseases in low resource settings. *International Journal of Computer Applications*, 187(14), 66–73. https://doi.org/10.5120/ijca2025925176

[8] Vinyals, O., Blundell, C., Lillicrap, T., Wierstra, D. (2016). Matching networks for one shot learning. Advances in neural information processing systems, 29. https://arxiv.org/abs/1606.04080

[9] Liu, S., You, Y., Su, H., Meng, G., Yang, W., Liu, F. (2022). Few-shot object detection in remote sensing image interpretation: Opportunities and challenges. Remote Sensing, 14(18), 4435. https://doi.org/10.3390/rs14184435

[10] Broni-Bediako, C., Xia, J., Song, J., Chen, H., Siam, M., Yokoya, N. (2024). Generalized few-shot semantic segmentation in remote sensing: Challenge and benchmark. IEEE Geoscience and Remote Sensing Letters, 21, 1-5. https://arxiv.org/abs/2409.11227

[11] Zhan, Y., Fu, K., Yan, M., Sun, X., Wang, H., Qiu, X. (2017). Change detection based on deep siamese convolutional network for optical aerial images. IEEE Geoscience and Remote Sensing Letters, 14(10), 1845-1849. https://doi.org/10.1109/LGRS.2017.2738149