
Mapping Farmed Landscapes from Remote Sensing

Michelangelo Conserva
Google Research

Alex Wilson
Google Research

Charlotte Stanton
Google Research

Vishal Batchu
Google Research

Varun Gulshan
Google Research

Abstract

To overcome the critical lack of detailed ecological maps needed for managing agricultural landscapes, we developed *Farmscapes*: the first large-scale, high-resolution map that identifies ecologically vital rural features, including often overlooked elements like hedgerows and stone walls. We achieved high accuracy in mapping key habitats with a deep learning model trained on aerial imagery and expert annotations. As a result, this work enables data-driven planning for habitat restoration, supports the monitoring of key initiatives like the EU Biodiversity Strategy, and lays a foundation for advanced analysis of landscape connectivity.

1 Introduction

Agricultural systems face increasing pressure to meet the growing demand for food while maintaining ecological balance and mitigating climate impacts (Webb et al., 2020). The intensification of agriculture over recent decades has led to significant environmental challenges, including habitat loss, soil degradation, and declining wildlife populations (Landis, 2017).

In intensively farmed landscapes, natural habitats such as woodlands, grasslands, and hedgerows are often fragmented or removed altogether (Krauss et al., 2010). In this scenario, remaining patches of semi-natural habitats are disconnected from one another, creating isolated islands in a sea of monocultures. This fragmentation limits species movement, reduces genetic exchange, and diminishes overall biodiversity (Wan et al., 2018). Biodiversity loss, in turn, weakens the resilience of ecosystems, reducing their ability to provide essential services such as pollination, pest control, and climate regulation services that are crucial to sustainable agriculture (Lanz et al., 2018).

The global urgency to address biodiversity loss is underscored by international commitments such as the Kunming-Montreal Global Biodiversity Framework, which includes the landmark “30 by 30” initiative to protect 30% of terrestrial and marine ecosystems by 2030 (UNEP-WCMC and IUCN, 2024; CBD, 2022). This ambitious goal has been adopted at a continental level by the EU Biodiversity Strategy for 2030, compelling Member States to enhance habitat protection and restoration, with a particular focus on farmland biodiversity (Hermoso et al., 2022).

While these initiatives reflect a growing consensus on the need for large-scale restoration, their success hinges on implementing targeted, cost-efficient interventions (Duke et al., 2013). A fundamental prerequisite for such strategic action is the availability of accurate maps of rural landscapes, which are essential for guiding governments and local initiatives in prioritising restoration efforts and allocating resources effectively (Turner, 2001; Moilanen et al., 2005).

Current efforts in high-resolution landscape mapping have shown promise but face significant limitations. While deep learning techniques are now state-of-the-art for segmenting remote sensing imagery, their application to fine-scaled rural features remains limited. Studies on hedgerow mapping often suffer from coarse image resolution that misses fine details, misclassification with other woody

vegetation, or methodologies that lead to imprecise boundaries (Ahlsvede et al., 2021; Strnad et al., 2023; Muro et al., 2024). Research on stone wall mapping has leveraged high-resolution LiDAR data, but the intensive manual labour required for labelling has restricted these efforts to smaller areas (Suh and Ouimet, 2023; Trotter et al., 2022). Furthermore, existing commercial and institutional datasets are often limited by high costs or a lack of methodological transparency, hindering widespread scientific use (Broughton et al., 2024; International, 2024). A critical gap therefore exists for a mapping resource that is simultaneously large-scale, high-resolution, and openly accessible.

To address these limitations, we introduce *Farmscapes*: a large-scale, high-resolution (25 cm) open dataset that identifies ecologically vital rural features across England, with preliminary, unvalidated maps generated for other European countries. This map was produced using a deep learning model trained on a novel dataset of landscape elements annotations collected for imagery in England. While quantitatively validated across most of England, the model has also been used to generate illustrative landscape maps for other European countries with similar rural characteristics. By releasing this extensive dataset publicly, we provide ecologists, policymakers, and local communities in England with a powerful, data-driven tool, while offering a foundation for a future pan-European resource. To address potential privacy concerns, the initial public release is limited to the hedgerow, stone wall, and woodland data layers. To the best of our knowledge, *Farmscapes* is the first resource of its kind, enabling precise, large-scale assessments of ecological status and facilitating the targeted restoration efforts needed to meet ambitious biodiversity goals.

2 Landscape elements

Agricultural landscapes are composed of a mosaic of distinct features, which we refer to as landscape elements. These elements, both natural and anthropogenic, define the physical structure of the environment and are fundamental to its ecological function. Their composition and spatial arrangement influence biodiversity, species movement, and the provision of ecosystem services. In this study, we focus on mapping four key elements that are vital to the character and health of European farmland: hedgerows, stone walls, woodland, and the surrounding farmed land.

Hedgerows are linear features composed of shrubs and trees that form boundaries in farmed landscapes. Their structure can vary, ranging from low, intensively managed shrubs to complex, unmanaged lines containing mature trees. This structural diversity is key to their ecological importance.

Woodland refers to areas dominated by trees with a more open canopy and lower density than forests, often featuring a developed understory. These ecosystems are crucial for carbon sequestration, soil stability, and water regulation, and they significantly enhance landscape diversity by acting as transitional zones between open land and dense forests.

Stone walls are man-made linear boundaries constructed from stone. While serving agricultural purposes like enclosing livestock, they also function as important microhabitats. The crevices provide shelter for insects, small animals, and reptiles, and support specialised plant life like mosses and lichens, thereby enhancing local biodiversity.

Farmed land comprises areas actively managed for agricultural production, including the cultivation of crops and the raising of livestock. As a dominant global land use, it is fundamental to food security and includes systems ranging from intensive monocultures to diverse agroforestry. In our study, this class represents the matrix in which other landscape elements are embedded.

3 Data

In order to train our deep learning model, we collect 25cm aerial imagery (input) and human annotations (labels). To ensure the accuracy of these annotations, we provided the annotators with 1m LiDAR-derived height maps as a supplementary reference tool.

Aerial imagery. Identifying fine-scaled features like hedgerows requires imagery with a resolution significantly higher than publicly available sources like Sentinel-2 (10m). To address this, we used proprietary 25cm resolution aerial imagery, captured over England (2018-2022). This high-resolution dataset served as the sole input for our deep learning model.

LiDAR measurements. While aerial imagery contains visual cues for height, such as shadows and texture, interpreting them consistently can be ambiguous. To establish a highly accurate and unambiguous ground truth, we equipped our annotators with a 1m resolution height map from the UK Environment Agency’s LiDAR Digital Terrain Model (DTM) (Agency, 2022). This reference layer allowed them to rapidly and confidently distinguish between features of different heights, such as a tall hedgerow and low scrub. The deep learning model was then trained on the aerial imagery to learn the correlation between these visual cues and the LiDAR-verified labels. Note that the LiDAR data was used exclusively for creating the ground truth and was not an input to the model itself.




Figure 1: Comparison of data sources.

Sampling Strategy. We created a novel dataset by sampling 942 image tiles from rural England, covering approximately 247 km². To ensure a representative sample of key features, our process first excluded high-density urban areas and large forests, then used stratified sampling to over-sample regions with a known high density of stone walls. Each tile measures 512m × 512m, which corresponds to 2048 × 2048 pixels at 25cm resolution.

Annotation Process. For annotation, experts used both the aerial imagery (Figure 1a) and LiDAR-derived height data (Figure 1b) as complementary references. This process resulted in a final annotation mask (Figure 1c) labeling seven exhaustive land cover classes: *farmed land*, *hedgerows*, *woodland*, *stone walls*, *other vegetation*, *water*, and *other* (encompassing roads, buildings, etc.).

Training targets. To model the landscape’s vertical structure, where hedgerows or woodland can overlay ground cover, we defined four training targets. A multiclass target was used for mutually exclusive ground classes (*farmed land*, *other vegetation*, *other*), while three separate binary targets were used to identify features that can co-exist within the same pixel space (*hedgerows*, *woodland*, *stone walls*). The distribution of training targets is shown in Table 1.

Table 1: Percentage of classes of the training targets.

	Ground	Hedgerow		Stone wall		Woodland	
Farmed land	67%	Non hedgerow	99%	Non stone wall	99%	Non woodland	89%
Vegetation	27%	Hedgerow	1%	Stone wall	1%	Woodland	11%
Other	3%						

Deep learning model. We use a SETR-PUP segmentation transformer (Zheng et al., 2021), which uses a Vision Transformer (ViT) encoder and separate convolutional decoders for each of the four targets. The encoder was pretrained on ~ 100 million global aerial images using a masked autoencoder approach (He et al., 2022). To improve generalisation, we applied extensive data augmentation, including random flips, rotations, color jitter, resizing, and cutout (DeVries, 2017). The model was trained on 512×512 pixel (128m × 128m) random crops.

4 Results

We evaluated our model’s performance on the hold-out test set using f1-score, precision, and recall. Classification thresholds were selected to maximise the f1-score on the validation set. F1-scores are presented in Table 2. Precision and recall are presented in the appendix (see Table 3). While these scores establish a strong baseline, direct quantitative benchmarking against state-of-the-art methods is precluded by our use of proprietary 25cm imagery, a step that was essential to resolve fine-scale features like stone walls.

The model demonstrates solid performance on broad landscape features. For major classes like **Woodland** and **Farmed land**, it achieved f1-scores of 96 ± 1 and 95 ± 1 , respectively. This indicates that the model reliably identifies the primary matrix of the agricultural landscape. The segmentation of fine-scale linear features presented a greater challenge. Nevertheless, the model achieves an f1-score of 72 ± 1 for **Hedgerows**. Performance on **Stone walls** was lower, with an f1-score of 60 ± 1 , reflecting the difficulty of identifying these features, which are often just a few pixels wide and can be obscured by vegetation. These results represent a significant advancement in mapping these critical, yet difficult-to-detect, ecological corridors.

Table 2: Test metrics with average and standard deviation over twelve seeds.

	Ground			Above ground		
	Vegetation	Farmed land	Other	Stone walls	Hedgerows	Woodland
f1-score	84 ± 1	95 ± 1	81 ± 3	60 ± 1	72 ± 1	96 ± 1

5 Conclusion

In this study, we have successfully demonstrated the power of combining high-resolution aerial imagery with a deep learning approach to produce *Farmscapes*: the first large-scale, open-access map of key rural landscape elements. Our model, which leverages a Vision Transformer architecture, has proven highly effective at segmenting the agricultural matrix, achieving excellent performance in mapping broad features like woodland and farmed land. More importantly, it marks a significant advancement in the automated identification of fine-scale linear features, delivering strong results for hedgerows and a promising baseline for the notoriously challenging task of mapping stone walls. By creating and publicly releasing this dataset, we provide a foundational tool for ecologists, land managers, and policymakers to support data-driven conservation, monitor progress towards ambitious goals like the EU Biodiversity Strategy, and enable new research into landscape connectivity.

This study is subject to four key limitations. First, our use of proprietary 25cm aerial imagery, while essential for resolving fine-scale features like hedgerows and stone walls, precludes direct quantitative benchmarking against state-of-the-art methods. Second, model performance is fundamentally constrained by annotation quality; despite our use of LiDAR to maximize accuracy, the manual labeling process is subject to human error, which can introduce noise. Third, this map represents a static snapshot from 2018-2022 imagery. Monitoring landscape change, a critical need for tracking policy effectiveness, is not yet possible and remains important future work. Finally, quantitative validation was conducted exclusively on data from England, meaning the map’s reliability outside this region remains unquantified; any application to other geographies must be considered preliminary.

Building on this foundation, future work will proceed along three main avenues. First, we aim to refine the model’s accuracy, challenging classes like stone walls, by exploring other architectures and complementary data sources. Second, we plan to broaden the map’s scope, geographically and temporally. This involves collecting new annotated data from diverse European landscapes to improve generalization and adapting the model to work with lower-resolution imagery to enable historical change analysis. Finally, our ultimate goal is to translate this detailed map into actionable ecological insights. By conducting a comprehensive network analysis of landscape connectivity, we can identify critical corridors, fragmentation points, and priority areas for restoration, transforming *Farmscapes* into a dynamic tool for strategic environmental planning.

References

- Agency, U. E. (2022). Uk environment agency's lidar digital terrainmodel (dtm). Accessed on 05.02.2014.
- Ahlswede, S., Asam, S., and Röder, A. (2021). Hedgerow object detection in very high-resolution satellite images using convolutional neural networks. *Journal of Applied Remote Sensing*, 15(1):018501–018501.
- Broughton, R., Burkmar, R., McCracken, M., Mitschunas, N., Norton, L., Pallett, D., Patton, J., Redhead, J., Staley, J., Wood, C., and Pywell, R. (2024). Ukceh land cover plus: Hedgerows 2016-2021 (england).
- CBD, U. (2022). Kunming-montreal global biodiversity framework. In *Fifteenth meeting of the Conference of the Parties to the Convention on Biological Diversity (Part Two) Decision 15/4*.
- Chen, L.-C. (2017). Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*.
- DeVries, T. (2017). Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*.
- Diakogiannis, F. I., Waldner, F., Caccetta, P., and Wu, C. (2020). Resunet-a: A deep learning framework for semantic segmentation of remotely sensed data. *ISPRS Journal of Photogrammetry and Remote Sensing*, 162:94–114.
- Dial, G., Bowen, H., Gerlach, F., Grodecki, J., and Oleszczuk, R. (2003). Ikonos satellite, imagery, and products. *Remote sensing of Environment*, 88(1-2):23–36.
- Duke, J. M., Dundas, S. J., and Messer, K. D. (2013). Cost-effective conservation planning: Lessons from economics. *Journal of environmental management*, 125:126–133.
- Hansen, M. C., Potapov, P. V., Moore, R., Hancher, M., Turubanova, S. A., Tyukavina, A., Thau, D., Stehman, S. V., Goetz, S. J., Loveland, T. R., et al. (2013). High-resolution global maps of 21st-century forest cover change. *science*, 342(6160):850–853.
- He, K., Chen, X., Xie, S., Li, Y., Dollár, P., and Girshick, R. (2022). Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16000–16009.
- Hermoso, V., Carvalho, S., Giakoumi, S., Goldsborough, D., Katsanevakis, S., Leontiou, S., Markantonatou, V., Rumes, B., Vogiatzakis, I., and Yates, K. (2022). The eu biodiversity strategy for 2030: Opportunities and challenges on the path towards biodiversity recovery. *Environmental Science & Policy*, 127:263–271.
- International, B. (2024). National hedgerow map.
- Krauss, J., Bommarco, R., Guardiola, M., Heikkinen, R. K., Helm, A., Kuussaari, M., Lindborg, R., Öckinger, E., Pärtel, M., Pino, J., et al. (2010). Habitat fragmentation causes immediate and time-delayed biodiversity loss at different trophic levels. *Ecology letters*, 13(5):597–605.
- Landis, D. A. (2017). Designing agricultural landscapes for biodiversity-based ecosystem services. *Basic and Applied Ecology*, 18:1–12.
- Lanz, B., Dietz, S., and Swanson, T. (2018). The expansion of modern agriculture and global biodiversity decline: an integrated assessment. *Ecological Economics*, 144:260–277.
- Marconcini, M., Metz-Marconcini, A., Üreyen, S., Palacios-Lopez, D., Hanke, W., Bachofer, F., Zeidler, J., Esch, T., and Strano, E. (2020). World Settlement Footprint (WSF) 2015.
- Moilanen, A., Franco, A. M., Early, R. I., Fox, R., Wintle, B., and Thomas, C. D. (2005). Prioritizing multiple-use landscapes for conservation: methods for large multi-species planning problems. *Proceedings of the Royal Society B: Biological Sciences*, 272(1575):1885–1891.

- Muro, J., Blickensdorfer, L., Don, A., Köber, A., Schwieder, M., and Erasmi, S. (2024). Hedgerow mapping with high resolution satellite imagery to support policy initiatives at national level. *Available at SSRN 5011709*.
- Neumann, J. L., Griffiths, G. H., Hoodless, A., and Holloway, G. J. (2016). The compositional and configurational heterogeneity of matrix habitats shape woodland carabid communities in wooded-agricultural landscapes. *Landscape Ecology*, 31:301–315.
- Pirbasti, M. A., McArdle, G., and Akbari, V. (2024). Hedgerows monitoring in remote sensing: A comprehensive review. *IEEE Access*.
- Strnad, D., Horvat, Š., Mongus, D., Ivajnsič, D., and Kohek, Š. (2023). Detection and monitoring of woody vegetation landscape features using periodic aerial photography. *Remote Sensing*, 15(11):2766.
- Suh, J. W. and Ouimet, W. (2023). Regional mapping of stone walls in northeastern us using deep learning. In *Geological Society of America Abstracts*, volume 55, page 386159.
- Trotter, E. F. L., Fernandes, A. C. M., Fibæk, C. S., and Keßler, C. (2022). Machine learning for automatic detection of historic stone walls using lidar data. *International Journal of Remote Sensing*, 43(6):2185–2211.
- Turner, M. (2001). *Landscape ecology in theory and practice*.
- UNEP-WCMC and IUCN (2024). *Protected planet report*.
- Wan, H. Y., Cushman, S. A., and Ganey, J. L. (2018). Habitat fragmentation reduces genetic diversity and connectivity of the mexican spotted owl: a simulation study using empirical resistance models. *Genes*, 9(8):403.
- Webb, P., Benton, T. G., Beddington, J., Flynn, D., Kelly, N. M., and Thomas, S. M. (2020). The urgency of food system transformation is now irrefutable. *Nature Food*, 1(10):584–585.
- Zheng, S., Lu, J., Zhao, H., Zhu, X., Luo, Z., Wang, Y., Fu, Y., Feng, J., Xiang, T., Torr, P. H., et al. (2021). Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6881–6890.

A Annotation methodology and guidelines

Landscapes elements. Figure 2 presents visual examples of the primary landscape elements targeted in this study: hedgerows, stone walls, and woodland.



Figure 2: Typical English landscape elements, hedgerows (left), stone walls (center), and woodland (right).

Hedgerow Annotation Guidelines. To ensure consistent labelling of hedgerows, which exhibit significant structural diversity, our annotation protocol follows the typology proposed by Neumann et al. (2016). This framework defines three distinct hedgerow types based on their structure and management intensity. The key rule for our annotation process was to map these types to our final land cover classes as follows:

- **Type 1:** A low-lying, intensively managed hedge composed of shrubs and lacking mature trees (Figure 3a). According to Neumann et al. (2016), these typically reach up to 1.5 meters in height with an average width of 2.5 meters.
- **Type 2:** A less managed, taller hedgerow containing small trees or tall shrubs (Figure 3b). These features are structurally more complex than Type 1, with an average width of 7 meters.
- **Type 3:** A hedgerow that includes mature trees, often appearing from an aerial perspective as a narrow, linear strip of woodland (Figure 3c). Due to its structural similarity to forest edges, this type was classified as *Woodland* to maintain a clear distinction based on the presence of a mature tree canopy.



(a) Hedgerow Type 1.

(b) Hedgerow Type 2.

(c) Hedgerow Type 3.

Figure 3: Types of hedgerows.

Sampling Strategy. The 942 image tiles for annotation were sampled from across England using a stratified approach designed to ensure representative coverage of agricultural landscapes while oversampling for the rare stone wall class. The process was as follows:

1. **Exclusion Masking:** We first excluded non-target areas. High-density urban regions were masked using the World Settlement Footprint (Marconcini et al., 2020), and large forest


blocks were masked using the Global Forest Change dataset (Hansen et al., 2013). This focused the sampling on the rural matrix.

2. **Stratification:** The remaining area was divided into two strata to ensure adequate representation of stone walls (Figure 4a). Stratum 1 consisted of regions with a known high density of stone walls (e.g. Cotswold and the Lake District), while Stratum 2 comprised all other eligible rural areas.
3. **Random Sampling:** We then performed stratified random sampling, allocating 15% of the total tiles to Stratum 1 and the remaining 85% to Stratum 2. Figure 4b shows the final distribution of sampled locations.

This strategy ensures the dataset captures a wide variety of rural landscapes while providing sufficient examples of the less common stone wall class for effective model training.



(a) High-density urban areas (white), forests (yellow), and areas with higher stone wall density (red).



(b) Locations sampled for annotation (red) superimposed on areas with available aerial imagery.

Figure 4: Sampling strategy illustration.

B Related works

In the context of high-resolution landscape mapping, hedgerow mapping is one of the most studied applications (Pirbasti et al., 2024). In Germany, Ahlswede et al. (2021) trained a DeepLab v3 (Chen, 2017) model using 1-meter resolution imagery from the IKONOS mission (Dial et al., 2003). Labels were manually digitized hedgerow polygons provided by the Bavarian State Office for the Environment. Although this study improved upon the use of coarser imagery, it was limited by the resolution, which may miss finer details such as small gaps in hedgerows critical for ecological analysis. Additionally, the focus solely on hedgerows excluded other important landscape features. Strnad et al. (2023) employed a U-Net model trained on aerial photography of Slovenia, with a high spatial resolution of 25 cm. Labels for this study were derived from a reference layer based on LiDAR point clouds collected in 2014 and processed to exclude buildings. While this approach allowed for scaling without manual labelling, it was limited to identifying woody vegetation and could not specifically distinguish hedgerows. More recently, Muro et al. (2024) used multi-temporal PlanetScope satellite imagery with a 3-meter resolution to map hedgerows across Germany using a U-Net architecture. Labels were sourced from a dataset created by the Schleswig-Holstein State Office for Agriculture, Environment, and Rural Areas, combining digital terrain models and high-resolution imagery. However, this approach buffered hedgerow labels by five meters, potentially leading to over-segmentation and the loss of fine details. In addition to academic research, institutional efforts have also played a role in advancing hedgerow mapping. The UK Centre for Ecology and Hydrology offers a dataset of linear hedgerows, providing a valuable resource for ecological studies (Broughton et al., 2024). Similarly, Bluesky (International, 2024) offers a map of hedgerows and trees across the UK, which incorporates detailed volumetric information. While these datasets offer significant

potential for large-scale analyses, their high cost and lack of transparency regarding methodologies and validation processes limit broader accessibility and utility.

In parallel, mapping stone walls has also benefited from advances in remote sensing and deep learning. Suh and Ouimet (2023) mapped stone walls in the Northeastern USA using U-Net, and Diakogiannis et al. (2020) models with high-resolution airborne LiDAR data. The model input consisted of LiDAR-derived hillshades and slope maps, with labels created through manual digitization of stone walls from LiDAR data, supplemented by aerial imagery, Google Street View, and field verification. Similarly, Trotter et al. (2022) focused on updating Denmark’s stone wall registry using a U-Net model applied to LiDAR-derived terrain data. The model input included a Digital Terrain Model (DTM), Height Above Terrain (HAT), and a Sobel-filtered DTM, all at a 40 cm resolution. Labels were generated from a stone wall dataset provided by the Danish Ministry of Culture, validated and corrected using the DTM. While both approaches demonstrate the effectiveness of using LiDAR data to detect stone walls, the high manual labor required to create the labels limits the scale of the datasets.

C Training details




Figure 5: Model architecture.

Hyperparameter optimization. The 942 tiles of the collected datasets are partitioned into 742 for training, 100 for validation, and 100 for testing. To accurately assess model performance, validation and test tiles were carefully hand-selected. This selection process prioritized tiles with the clearest annotations and a representative number of instances for rarer classes like stone walls and hedgerows. We optimized the training of hyperparameters by sweeping over weight decay and learning rate. Performance for each hyperparameter combination was averaged across three random seeds evaluated on the validation set.

Table 3: Test metrics with average and standard deviation over twelve seeds.

	Ground			Above ground		
	Vegetation	Farmed land	Other	Stone walls	Hedgerows	Woodland
f1-score	84 ± 1	95 ± 1	81 ± 3	60 ± 1	72 ± 1	96 ± 1
recall	90 ± 1	91 ± 1	90 ± 1	53 ± 1	70 ± 1	92 ± 1
precision	85 ± 1	96 ± 1	82 ± 4	58 ± 1	73 ± 1	90 ± 1

D Detailed performance metrics

Ground Layer. For the *Vegetation* class, the model achieves an f1-score of 84 ± 1 , a recall of 90 ± 1 , and a precision of 85 ± 1 . This indicates that the model can accurately identify general vegetation cover with high consistency and completeness. The *Farmed land* class exhibits even higher performance, with an f1-score of 95 ± 1 , a recall of 91 ± 1 , and a precision of 96 ± 1 , demonstrating the model’s strong ability to delineate agricultural areas. The *Other* class, encompassing a variety of non-vegetated, non-farmed land covers, achieves an f1-score of 81 ± 3 , a recall of 90 ± 1 , and a precision of 82 ± 4 . While the precision shows some variability, the high recall indicates that the model captures most of this diverse class.

Stone Walls. The model achieves an f1-score of 60 ± 1 for stone walls. The recall is 53 ± 1 , and the precision is 58 ± 1 . This suggests that while the model has some ability to identify stone walls, there’s room for improvement, particularly in capturing all instances (recall). The relatively balanced precision and recall indicate that the model is neither overly prone to false positives nor false negatives, but rather struggles with consistent detection of these narrow features.

Hedgerows. For hedgerows, the model exhibits an f1-score of 72 ± 1 , a recall of 70 ± 1 , and a precision of 73 ± 1 . These results demonstrate a good ability to identify hedgerows, with a reasonable balance between correctly identifying them (precision) and capturing all instances (recall). The performance on hedgerows is notably higher than that on stone walls, likely due to their typically larger size and more distinct visual characteristics.

Woodland. The model performs exceptionally well on the *Woodland* class, achieving an f1-score of 96 ± 1 , a recall of 92 ± 1 , and a precision of 90 ± 1 . This indicates that the model can very accurately and consistently identify woodland areas. The high f1-score, combined with high recall and precision, demonstrates the model’s strong capability in delineating this important land cover type.

E Qualitative examples



Figure 6: English landscape.



Figure 7: France landscape.



Figure 8: Germany landscape.




Figure 9: Sweden landscape.




Figure 10: Italy landscape.




Figure 11: Portugal landscape.