



# Habitat classification from ground-level imagery using deep neural networks

Hongrui Shi <sup>a,1b,\*</sup>, Lisa Norton <sup>b</sup>, Lucy Ridding <sup>b</sup>, Simon Rolph <sup>c</sup>, Tom August <sup>c</sup>, Claire M. Wood <sup>b</sup>, Lan Qie <sup>d</sup>, Petra Bosilj <sup>a,e</sup>, James M. Brown <sup>a</sup>

<sup>a</sup> School of Engineering & Physical Sciences, University of Lincoln, Brayford Pool, Lincoln, LN6 7TS, United Kingdom

<sup>b</sup> UK Centre for Ecology & Hydrology, Lancaster Environment Centre, Library Avenue, Bailrigg, Lancaster, LA1 4AP, United Kingdom

<sup>c</sup> UK Centre for Ecology & Hydrology, Maclean Building, Benson Lane, Crowmarsh Gifford, Wallingford, Oxfordshire, OX10 8BB, United Kingdom

<sup>d</sup> School of Natural Sciences, University of Lincoln, Brayford Pool, Lincoln, LN6 7TS, United Kingdom

<sup>e</sup> Department of Advanced Computing Sciences, Maastricht University, Maastricht, 6229 GS, The Netherlands

## ARTICLE INFO

Dataset link: <https://github.com/WhiteGiveFive/Habitat-Classification-from-Ground-Level-Imagery>

### Keywords:

Habitat classification  
Environmental monitoring  
Deep learning  
Vision transformers  
Conservation  
Biodiversity  
Ground-level imagery

## ABSTRACT

Habitat assessment at local scales — critical for enhancing biodiversity and guiding conservation priorities — often relies on expert field surveys that can be costly, motivating the exploration of AI-driven tools to automate and refine this process. While most AI-driven habitat mapping depends on remote sensing, it is often constrained by sensor availability, weather, and coarse resolution. In contrast, ground-level imagery captures essential structural and compositional cues invisible from above and remains underexplored for robust, fine-grained habitat classification. This study addresses this gap by applying state-of-the-art deep neural network architectures to ground-level habitat imagery. Leveraging data from the UK Countryside Survey covering 18 broad habitat types, we evaluate two families of models — convolutional neural networks (CNNs) and vision transformers (ViTs) — under both supervised and supervised contrastive learning paradigms. Our results demonstrate that ViTs consistently outperform state-of-the-art CNN baselines on key classification metrics (Top-3 accuracy = 91%, MCC = 0.66) and offer more interpretable scene understanding tailored to ground-level images. Moreover, supervised contrastive learning significantly reduces misclassification rates among visually similar habitats (e.g., Improved vs. Neutral Grassland), driven by a more discriminative embedding space. Finally, our best model performs on par with experienced ecological experts in habitat classification from images, underscoring the promise of expert-level automated assessment. By integrating advanced AI with ecological expertise, this research establishes a scalable, cost-effective framework for ground-level habitat monitoring to accelerate biodiversity conservation and inform land-use decisions at a national scale.

## 1. Introduction

Habitats are generally classified according to underlying biophysical characteristics of the land on which they occur and the associated species assemblages. Accurate habitat classification is essential for biodiversity conservation and effective land management. Habitat extent and condition data underpin policy frameworks—such as UK's Biodiversity Net Gain (zu Ermgassen et al., 2021; GOV.UK, 2023) and EU Biodiversity Strategy for 2030 (European Commission, 2020)—and guide conservation priorities, yet they currently depend on expert field surveys which can be costly or difficult to access, for example for landowners looking to understand their land and its value in relation to offsetting or carbon credits (Schofield et al., 2024). Recent advances in artificial intelligence (AI) have demonstrated strong potential to automate image-based ecological classification tasks, making habitat

identification more accessible for non-experts, thereby reducing costs and alleviating professional workloads.

To date, most AI-driven habitat mapping works rely on remote sensing data (e.g., satellite or drone imagery) (Van An et al., 2023; Norris et al., 2024). Though remote-sensing-based habitat classification excels at quantifying spatial extent, it can be hindered by sensor availability, weather conditions, and relatively coarse spatial resolution (satellite imagery typically several m/pixel) (Díaz-Ireland et al., 2024; Morueta-Holme et al., 2024; Marjani et al., 2025). Ground-level imagery, by contrast, captures local structural and compositional cues for habitats (Perrett et al., 2023), which are invisible from above but critical for precise habitat classification. It is effective for site-scale habitat assessment and species identification and can be collected with consumer devices rather than a complex remote-sensing pipeline. However, this approach has been less studied for AI-driven habitat

\* Corresponding author.

E-mail address: [hshi@lincoln.ac.uk](mailto:hshi@lincoln.ac.uk) (H. Shi).

<https://doi.org/10.1016/j.ecoinf.2026.103751>

Received 18 July 2025; Received in revised form 26 March 2026; Accepted 28 March 2026

Available online 12 April 2026

1574-9541/© 2026 Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

classification. In the UK, image analysis is widely used to help both expert and non-expert field surveyors identify species, but no such equivalent currently exists for habitats. Yet, habitat classification is important for mapping and decision-making regarding land use and management, as well as in the context of BNG, a government supported approach currently in use in England to ensure that new projects deliver measurable improvement in nature either on site, or in areas away from the development site (offsets) (Mall et al., 2023; Morueta-Holme et al., 2024). In this context, we approach automatic habitat classification using ground-level imagery. By leveraging state-of-the-art AI techniques on richly annotated images from the UK Countryside Survey (Carey et al., 2008), we seek to support cost-effective site-level habitat monitoring, land cover mapping, and complement established approaches such as remote sensing.

Recently, data-driven computer vision has become instrumental in ecological and environmental research by leveraging the growing volume of imagery from diverse sensors to deliver social and economic benefits (Reynolds et al., 2025; Oba and Doi, 2025). Applications include species detection and monitoring, where deep learning models automate identification and population assessment from camera traps (Ratnayake et al., 2021; Bjerge et al., 2022; Xu et al., 2024); underwater and benthic habitat surveys (Zhong et al., 2023), which leverage images recorded by underwater cameras to monitor coral reef growth; plant recognition (Joly et al., 2024; Wang, 2025), employing large public datasets to train deep neural networks to classify plant species (e.g. PlantCEF (Pl@ntNet Consortium, 2025)); and habitat classification (Van An et al., 2023; Norris et al., 2024).

Within these applications, habitat classification is most directly aligned with our work. We therefore review it in more detail, focusing on data sources and model architectures where research gaps emerge. Currently, habitat classification remains particularly reliant on remote sensing, including satellite data (Praticò et al., 2021; Sittaro et al., 2022; Marcinkowska-Ochtyra et al., 2023; Van An et al., 2023) and drone-based images (Retallack et al., 2022; Norris et al., 2024). While these sources enable coverage of large areas, they depend on favourable environmental conditions, require specialised devices and expertise, and focus on a single top-down perspective, which limits their practicality for non-experts seeking timely, user-friendly, and precise habitat evaluations (Morueta-Holme et al., 2024). Compared to satellite imagery, ground-level imagery can be easily captured by on-site surveyors or even citizen scientists using consumer devices, offering a cost-effective alternative for detailed habitat classification. For example, DeepVerge (Perrett et al., 2023) demonstrates that street-view images can be used to train CNNs for counting grassland species. Saadeldin et al. (2022) shows that a simple CNN trained on a limited number of ground-level imagery achieves accurate classifications over three categories of grassland management intensity. Martinez-Sanchez et al. (2024) applies semantic segmentation to decompose visual features in ground-level photos and uses these features to classify coarse habitat classes (e.g., grassland, cropland), equivalent to UKHab Level 2 categories. However, despite this potential, ground-level imagery remains underexplored to address habitat classification across a variety of subcategories (e.g. different grasslands, woodlands, wetlands).

Regarding the adopted models for habitat classification, previous studies (Gómez-Ríos et al., 2019; Praticò et al., 2021; Sittaro et al., 2022; Marcinkowska-Ochtyra et al., 2023; Van An et al., 2023) have predominantly used classical machine learning models, such as support vector machines (Cortes and Vapnik, 1995) and random forests (Breiman, 2001), as well as deep neural networks, specifically CNNs, such as VGG (Simonyan and Zisserman, 2014) and ResNet (He et al., 2016). While these models are effective, they can struggle to address the increasing challenges of learning from more complex images and the growing demand for generalising to a wide range of habitat types. Conversely, the Vision transformer (ViT) (Dosovitskiy et al., 2020) has emerged as a state-of-the-art alternative to CNNs in various computer vision tasks. Compared to the CNN that only extracts

isolated regional features, the self-attention mechanism (Vaswani et al., 2017)—the foundation of ViT—allows this model to capture relationships between visual features across the entire image, leading to superior performance and interpretability in many applications. Recent habitat-focused studies incorporating ViTs (Jamali et al., 2023; Diaz-Ireland et al., 2024; Marjani et al., 2025) have shown promising results, but remain limited to a narrow set of habitats (e.g., grassland, wetland) and typically use satellite imagery rather than ground-level photographs. Consequently, there is a clear opportunity to adapt the ViT to more comprehensive, ground-level habitat classification.

This work aims to develop a model that is capable of automatically classifying diverse habitats based on ground-level imagery and with the long-term vision of deploying it on end-user devices in real-world settings (deployment itself falls beyond the scope of this paper). We leverage deep learning approaches, which have become dominant in various computer vision applications, to achieve this goal. A deep neural network is trained on thousands of labelled ground-level images provided by UK Centre for Ecology & Hydrology (UKCEH) Countryside Survey (CS) (Carey et al., 2008), which represents a wide variety of environments across the UK. Based on the UK Habitat (UKHab) Classification (UKHab Ltd, 2023) with a hierarchical habitat system, each image is annotated by experts with one of the level 3 (L3) habitat subcategories—fine-grained subdivisions of level 2 (L2) habitats (for example, grassland at L2 is split into types such as Acid Grassland, Improved Grassland, etc.). Appendix C provides detailed descriptions of L2 and L3 habitats as defined in the UKHab.

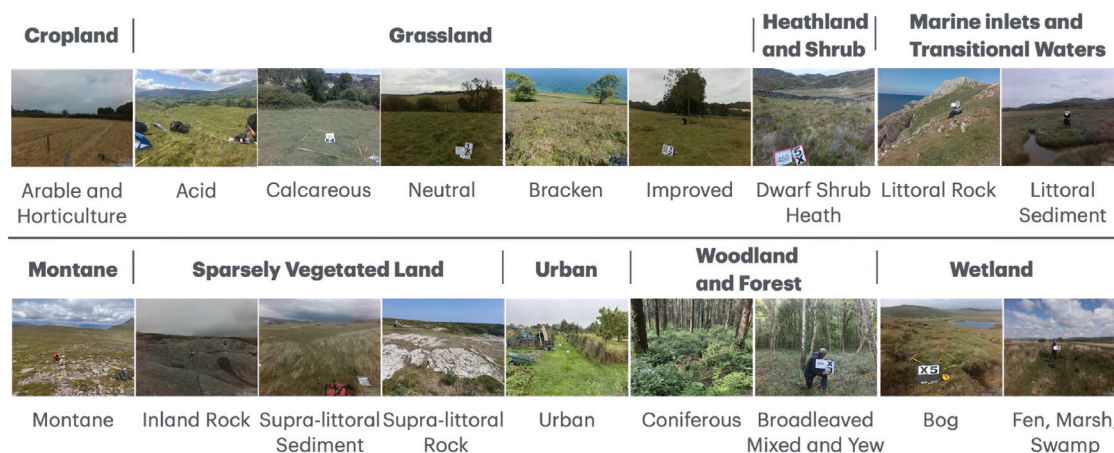
Two challenges relevant to ground-level habitat images are identified in the model development: (1) The model needs to classify habitats based on broad, open scenes with the presence of potentially distracting image features. Ground-level habitat images typically display large, continuous regions (e.g., swaths of grassland) that require an expansive, context-aware understanding of the scene. (2) Some habitat types (for example, Neutral Grassland and Improved Grassland) are visually similar, sharing subtle visual characteristics that makes them difficult to discern. Our training approach must learn distinct feature representations for such habitats so the classifier can differentiate them accurately. Fig. 1 displays habitat samples in the CS dataset, as well as the visual challenges of developing the model for automatic habitat classification.

To explore how best to meet these challenges, we structured our methodology around two axes of exploration:

- **Model architectures.** Convolutional Neural Networks (CNNs) are adopted as our baseline, reflecting the prevailing deep learning model in image classification. We further study Vision Transformer (ViT) (Dosovitskiy et al., 2020), whose self-attention mechanism (Vaswani et al., 2017) can more naturally capture global context in ground-level photographs.
- **Training paradigms.** Supervised learning with a standard cross-entropy objective is first employed for the multiclass habitat classification problem, to establish performance baselines for each architecture. We then investigate Supervised Contrastive Learning (SupCon) (Khosla et al., 2020) to pretrain the image encoder, aiming to learn tighter, more separable representations for visually similar habitats.

In summary, our research makes four main contributions:

- **Ground-level habitat classification benchmark.** Our benchmark leverages ground-level imagery to address a gap in tools requiring habitat classification (including carbon footprinting, biodiversity offsetting, etc.) for non-experts. It expands the ecological impact of AI by providing more accessible habitat assessments, allowing for extensive habitat classification across the country.
- **Model architecture finding.** Our systematic comparison of representative CNNs and ViTs demonstrates that ViTs consistently outperform comparable CNNs under a broad range of evaluation metrics for habitat classification.



**Fig. 1.** Examples of level 3 (L3) habitats defined by UKHab in the Countryside Survey dataset, grouped by their coarse L2 categories (bold text). Some L2 habitats, such as Cropland, only have one L3 class. Note that Improved Grassland, Montane, and Bracken have different origins than L3 habitats in UKHab but are treated as L3 habitats in the CS dataset. An explanation is provided in Section 2.1.

- Training paradigm investigation. We investigate a two-stage supervised contrastive (SupCon) learning paradigm (contrastive pretraining and classifier fine-tuning), which delivers modest gains over classical supervised learning for our best-performing model and reduces major confusions (e.g., Improved and Neutral Grassland), thereby improving reliability for real-world deployment.
- Human expert comparison. We compare our best model with three experienced ecologists. On a stratified subset, ViT combining with SupCon is competitive with the best human expert in interpreting habitats from level 3 habitat photographs, while revealing limits and opportunities for our approach.

## 2. Materials and methods

### 2.1. Data

Our study focuses on the habitats of the United Kingdom. The habitat images come from the UK Countryside Survey (CS) which covers a randomly stratified sample of UK countryside (Wood et al., 2017), meaning that habitats are sampled in proportion to their occurrence in the wider countryside. Images used were from the most recent rolling survey (CS2019-2023), which covered 500 1 km squares across 5 years, during which between 10 and 16 individuals (in each year) were employed as field surveyors. Around half of field surveyors are permanent staff members at UKCEH with other expert surveyors brought in as required. Field surveyors record species within vegetation plots between May and September of the survey year and take photographs of the plots within the context of the habitat in which they are placed, primarily for plot relocation purposes. Specifically, the surveyor is advised to record the location of the plot by taking several ground-level photographs at a near-earth height (1–2 m above ground) using smartphones, including plot markers in the foreground, the wider context of the plot and any significant (easily refindable and preferably permanent features) in the background where possible. These photographs depict the habitat context in the plot. Plots are assigned to habitats by the field surveyors according to the Broad and Priority Habitat classifications (U.K. Biodiversity Group, 2008). Broad Habitats comprise part of the hierarchical category system described in the UK Habitat Classification (UKHab) (UKHab Ltd, 2023).

UKHab defines five habitat levels, numbered 1 to 5, with higher levels (e.g., Level 3) representing more fine-grained habitat classifications than lower ones (e.g., Level 2). In this study, we particularly focus on classification at level 3 (L3) as level 1 (L1) and level 2 (L2) classifications are too coarse to add practical values for automatic

habitat classification and level (L4) was considered too fine-grained. However, L2 classes are still used in this work to describe grouped performance on L3. UKHab is one of several habitat classifications used in the UK but was chosen because it provides a comprehensive unified approach for all UK countries and has been developed to ensure consistency in classification over time and space which is suitable for digital mapping.

Initially, the data was cleaned by removing unlabelled habitat images, resulting in a dataset with a total number of 5598 samples, including 17 Broad L3 classes and one L4 class (Bracken), corresponding to 9 coarse L2 classes, as shown in Fig. 1. While Bracken is technically defined as an L4 habitat in UKHab, it is also classified as a Broad Habitat and is labelled by field surveyors in the CS dataset alongside the L3 habitats. Improved Grassland and Montane are originally Broad Habitats defined in the UK BAP Priority Habitats system (Joint Nature Conservation Committee, 2025), but they correspond to the scope of L3 habitats in UKHab and are utilised by field surveyors in the CS dataset. For simplicity and consistency throughout this study, we treat Bracken, Improved Grassland, and Montane as equivalent to L3 classes in UKHab, and all references to 'L3 habitats' in the CS dataset include these three classes unless otherwise specified.

Fig. 2 further depicts the distributions of the habitats at L3 and L2 levels. It shows that the CS dataset has a highly-tailed distribution. The grassland, cropland, wetland, woodland, and heathland are the most common habitats, constituting 98% of the dataset. We will refer to these habitats as *major habitats*.

### 2.2. Models

The backbone of our habitat classification system is a deep neural network capable of extracting discriminative features from ground-based imagery. In this work, we compare two architectures that have achieved state-of-the-art performance in image classification: conventional convolutional neural networks (CNNs) and Vision Transformers (ViTs). On our CS dataset, each model is trained to map raw image input to habitat labels by learning representations that capture visual patterns in diverse environments.

#### 2.2.1. Model architecture

**Convolutional neural networks.** The convolutional operation is the core of a CNN for modelling images. In particular, convolutional filters represented by matrices slide across the input image and perform element-wise multiplication and summation to detect visual features, such as edges, shapes, and colours (Goodfellow et al., 2016). The convolutional operation processes the images in local patches. Each

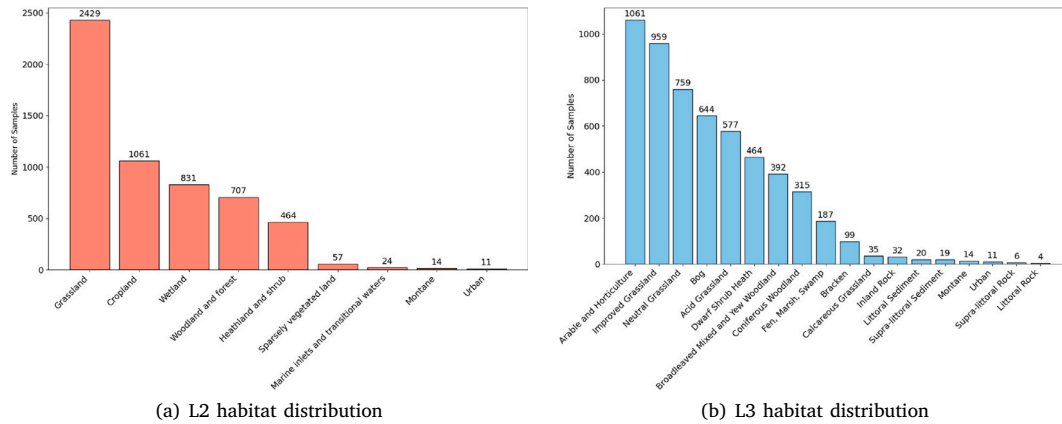


Fig. 2. Habitat distributions (L2 and L3) in the CS dataset based on the UKHab system.

time it is calculated, only a small region on the image is activated, whereas the remaining areas are not taken into account. This locality bias assumes that spatially proximate pixels are highly correlated, enabling efficient extraction of meaningful visual features. However, in ground-level habitat imagery—where scenes often contain broad scenes and multiple, potentially distracting objects—this local focus may limit the ability of a CNN to match interpretable visual features to the correct label.

**Vision transformers.** In contrast, ViT is capable of capturing long-range contextual relationships in the image with the attention mechanism. Specifically, ViT decomposes the input image into a sequence of patches. Each patch represents a small region of the input image and is encoded by the model into a compact representation. To model the relationship of the patches, the attention computes pairwise interactions among all patch representations as follows:

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d}}\right)\mathbf{V} \quad (1)$$

Here,  $\mathbf{Q}$ ,  $\mathbf{K}$ , and  $\mathbf{V}$  are the query, key matrix, and value matrices obtained from linear projections of the patch representations. The matrix product  $\mathbf{Q}\mathbf{K}^T$  essentially calculates pairwise inner products among projected representations, measuring their similarity. A higher similarity between two patch representations indicates they have a stronger connection. In this way, a broad scene with uniform visual features—such as grasses, woods, or crops—across the entire image can be connected together to represent the habitat. We hypothesise that the global perception of ViT is particularly advantageous for our habitat classification task: compared to the locality bias of CNNs, the non-local interaction enables ViT to learn image representation that encodes long-range regional dependencies and the overall contextual information of the scene, which are critical for generalising on ground-level imagery and interpreting the model performance.

Fig. 3 illustrates the differences of CNN and ViT on modelling images using an exemplary woodland image from the CS dataset.

### 2.2.2. Studied models

We selected representative CNN and ViT models that balance performance and efficiency for the architecture comparison. Using a size sweep to cover a range of capacities, we compared 9 models in total. With respect to the ViT architecture, we employed the Swin Transformer (SwinT) (Liu et al., 2021) for its computational efficiency. Compared to the basic ViT (Dosovitskiy et al., 2020), SwinT proposes the shifted window scheme to significantly reduce the computational workload of self-attention, while achieving classification performance on par with or superior to basic ViTs. SwinT has become a standard

representative ViT baseline for benchmarking in a number of state-of-the-art computer vision studies (Liu et al., 2022; Zhu et al., 2023; Wang et al., 2023). The efficiency and capability of SwinT enable us to accelerate the prototyping.

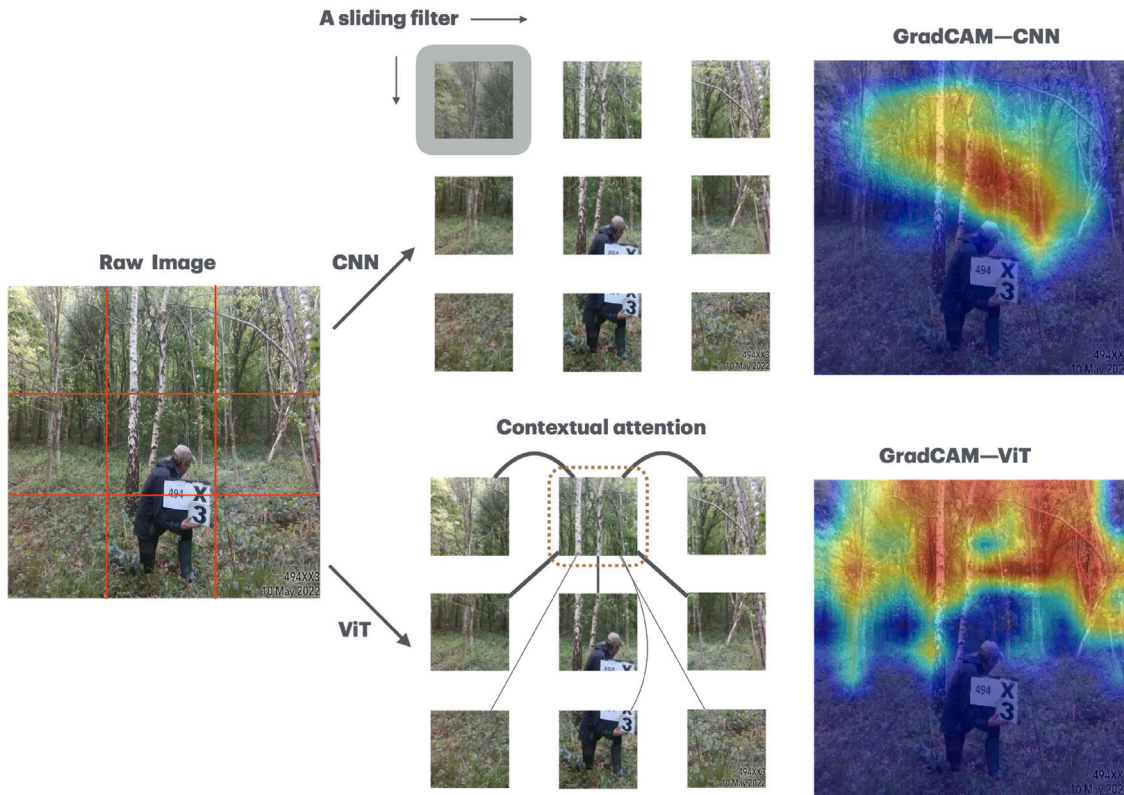
For the CNN baselines, we utilised three strong CNN models with complementary design principles, each of which has consistently ranked among the top performers on public image classification benchmarks such as CIFAR (Krizhevsky et al., 2009) and ImageNet (Deng et al., 2009). They are (1): Wide Residual Network (WRN) (Zagoruyko and Komodakis, 2016), which increases the number of filters to improve the performance of the popular ResNet (He et al., 2016); (2) EfficientNetV2 (Tan and Le, 2021), developed by neural architecture search (NAS) and offering better speed–accuracy trade-offs on ImageNet than both larger CNNs and the basic ViT; and (3) ResNeXt (Xie et al., 2017) introduces cardinality—the number of parallel paths in each block—as a new dimension for scaling CNNs, and achieves superior performance to ResNet on image classification benchmarks.

Additionally, we varied the sizes of the chosen models to provide a more comprehensive comparison, as well as guide our future model deployment on user-end devices, where lightweight models with lower resource demands and competitive capabilities will be favoured. Specifically, SwinT includes size-based variants of small, base, and large models, denoted as SwinT-S, SwinT-B, SwinT-L. In terms of CNNs, WRN-50-2, WRN-101-2, EfficientNetV2-M, EfficientNetV2-L, ResNeXt-50, ResNeXt-101 are chosen for the CNN baselines. Our choices of the ViT and CNN models are not exhaustive, but they are all proven models and can serve well as representatives for their family, enabling robust studies for our habitat classification task.

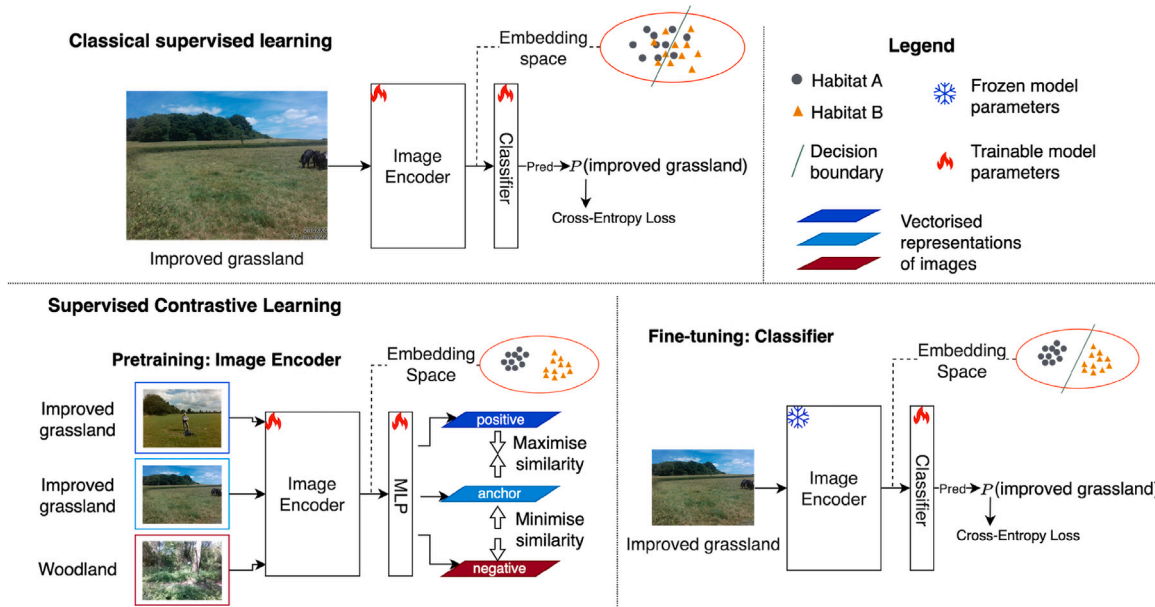
We followed the widely-accepted pretraining practice to avoid training the model from scratch. All of the chosen models are pretrained on the ImageNet dataset (Deng et al., 2009) and open-weighted. Our pretrained CNNs and SwinT are downloaded from the TorchVision library and the model repository of Hugging Face respectively.

### 2.3. Training methods

As the CS dataset is labelled by experts in the field (ground-truthed), we begin by fitting the chosen model via classical supervised learning: minimising the cross-entropy loss between the model’s predictions and the ground-truth habitat labels. However, visually similar habitat types pose significant challenges to model optimisation under the classic supervised learning paradigm. Illustrated in Fig. 4(a), some habitat classes—such as Neutral Grassland vs. Improved Grassland—share very similar visual patterns that are too subtle to distinguish. When extracting the embeddings of samples from the outputs of the image encoder, which is just ahead of the classifier, and project them into a 2D space, we anticipate to observe significant overlaps between classes (shown as



**Fig. 3.** The conceptual difference between how CNN and ViT process the same ground-level habitat image. A ground-level photograph of a mixed woodland scene is partitioned into a conceptual  $3 \times 3$  grid of patches. In the top row, the CNN extracts visual features using a sliding filter that only sees one patch at a time, making it tend to focus on prominent features in local regions. This locality is reflected in the GradCAM heatmap, where the model concentrates on the trees in the central region (heated area). In contrast, the ViT (bottom row) allows each patch to “attend” to every other patch through self-attention links. The lines in the diagram are a schematic way to show this idea—the thicker lines simply indicate stronger relationships between patches. The ViT therefore considers the global context of the scene, linking near and far regions. The GradCAM heatmap of ViT highlights areas not only on the central trees but also those along the edges, reflecting a broader contextual understanding. This behaviour aligns more closely with human interpretation of ground-level habitat photos, where habitat context often spans the entire scene rather than a single region.



**Fig. 4.** Comparison of Classical Supervised Learning and Supervised Contrastive Learning. Top (Classical Supervised Learning). An image encoder and classifier are trained end-to-end with cross-entropy. Because embeddings from different classes remain entangled, the classifier struggles to carve out a decision boundary in a poorly separated embedding space, which can hinder performance. Bottom (Supervised Contrastive Learning). First, the encoder is pretrained with a contrastive loss that pulls together same-class embeddings and pushes apart different-class embeddings. This step learns an encoder that generates well-clustered representations. A classifier is then fine-tuned on these fixed embeddings using classic supervised learning, enabling it to learn a much clearer decision boundary.

a conceptual illustration in Fig. 4). This low inter-class variance in the embedding space makes it difficult to learn a classifier that can carve out a clean decision boundary through the conventional supervised approach.

To encourage more separable embeddings, we use contrastive learning (Wu et al., 2018) for a comparison with the conventional supervised learning. Built on the prior knowledge that a model should learn invariant representations for similar samples, contrastive learning pulls together any given training image, denoted as an *anchor* and its augmented views (*positive samples*), while pushing apart all other different images (*negative samples*) in the embedding space. Contrastive learning enables the model to learn meaningful representations for images without references to ground-truth labels.

When labels are available, contrastive learning can be more flexible for learning representations. Particularly, Supervised Contrastive Learning (SupCon) (Khosla et al., 2020) integrates the label information into the contrastive loss to force the model to learn close embeddings for instances belonging to the same class, whereas embeddings of different classes are pushed away. As such, SupCon explicitly decreases intra-class variance and boosts inter-class variance, creating tight clusters for each habitat type and clear margins between them. With more separable embeddings, our hypothesis is that SupCon can alleviate the difficulties of learning an effective classifier to discern visually similar habitats.

To implement SupCon, we followed the two-stage transfer-learning strategy proposed in the original SupCon paper: (1) Pretraining: a Multi-Layer Perceptron (MLP) was attached to the image encoder and they were trained together with the SupCon loss on the CS training set. This phase aimed to train the image encoder to represent images into a space where same-class samples cluster and different-class samples repel. (2) Transfer learning: we discarded the MLP, froze the pretrained image encoder, and appended a fresh linear classifier to the image encoder. This classifier was then trained on the CS training set by the standard supervised learning with cross-entropy loss. At inference, we used the image encoder and the classifier trained in the transfer learning stage to predict habitat classes. Fig. 4 illustrates the training pipelines of classic supervised learning and supervised contrastive learning.

## 2.4. Evaluation

Our models are evaluated with respect to their classification performance and interpretability in predictions, as well as benchmarked against domain specialists.

### 2.4.1. Classification performance

We employ Top-1 accuracy, Top-3 accuracy, and Matthews Correlation Coefficient (MCC) to evaluate the classification performance on the test set. They are defined as follows:

$$Accuracy@1 = \frac{c^{(1)}}{s} \quad (2)$$

$$Accuracy@3 = \frac{c^{(3)}}{s} \quad (3)$$

$$MCC = \frac{c^{(1)}s - \sum_k t_k p_k}{\sqrt{s^2 - \sum_k p_k^2} \sqrt{s^2 - \sum_k t_k^2}} \quad (4)$$

where  $c^{(1)}$  is the total number of samples correctly predicted by the top-1 prediction,  $c^{(3)}$  is the total number of samples correctly predicted by the top-3 predictions,  $s$  is the total number of samples,  $t_k$  is the number of times class  $k$  truly occurred, and  $p_k$  is the number of times class  $k$  was predicted (top-1).

Top-1 accuracy is an evaluation metric that is broadly used to show the percentage of the instances that are correctly classified. Top-3 accuracy is used from a more practical perspective of habitat classification, where the model can still be helpful by producing the top three most

possible labels rather than just one. MCC is employed due to the class imbalance of the CS dataset, capturing the model performance while taking into account the highly skewed nature of this dataset (see Fig. 2). Additionally, we use Precision, Recall, and  $F_1$  score to evaluate individual habitat performance. Class-wise, they are defined as follows:

$$Precision = \frac{TP}{TP + FP} \quad (5)$$

$$Recall = \frac{TP}{TP + FN} \quad (6)$$

$$F_1 = \frac{2TP}{2TP + FP + FN} \quad (7)$$

where TP is the number of true positives, FN is the number of false negatives, and FP is the number of false positives. We further calculate the weighted  $F_1$  score for the overall classification performance while accounting for class imbalance. Weighted  $F_1$  calculates the  $F_1$  score for each class, then takes a support-weighted average of these per-class values, where ‘‘support’’ is the number of true instances of each class. Moreover, confusion matrix (CM) normalised by the ground-truth labels (columns add to one) is also utilised to illustrate classification performance for each habitat.

### 2.4.2. Model interpretability

GradCAM (Selvaraju et al., 2017) is utilised for explaining model predictions. GradCAM is a widely accepted technique to generate visual explanations for predictions made by deep neural networks on vision tasks. For image classification tasks, GradCAM produces a class activation map (CAM) using the gradients of the target class to highlight the crucial regions in the image for the model to predict the class. The heated regions shown in warm colours on the map suggests the class is determined by the visual features within these areas, and vice versa. GradCAM can be used to visually examine if model predictions are close to human perceptions. For instance, the GradCAM of a grassland image is expected to highlight regions that are largely occupied by grasses rather than the field surveyor.

### 2.4.3. Evaluation against domain specialists

We employ domain specialists to perform a blind annotation exercise on a subset of the test set, aiming to reveal the performance gaps between our model and human experts in evaluating habitat photographs. Specifically, we randomly selected 10% of the test set, 158 samples in total, for expert review. The draw was stratified to ensure that the subset preserved the distribution of the original test set. Our human specialist benchmark consists of (i) a research scientist with a broad international background and familiarity with UK conservation management, (ii) a research scientist with expertise in long-term land cover dynamics, drivers of vegetation change, habitat management and restoration, and (iii) a senior landscape ecologist who has been working on farmland biodiversity and the UK Countryside Survey for many years. We evaluate each expert’s annotations against the ground-truth labels using Top-1 accuracy, Top-3 accuracy, weighted  $F_1$  score, and MCC.

## 2.5. Implementation details

We constructed the training and test sets by splitting the raw CS dataset into 4200 training instances (75% of CS dataset) and 1398 test instances (25% of CS dataset). Applied with stratified partition, the resulting sets preserve approximately the same class distribution as the complete CS dataset. As the habitats of Supra-littoral Rock and Littoral Rock have too few samples in the CS dataset, the stratified partition did not include them in the test set. Additionally, 20% of instances in the training set were used for validating hyperparameter choices with the grid search strategy.

Our implementation was based on the PyTorch library. The AdamW (Loshchilov and Hutter, 2017) optimiser was chosen for model training.

**Table 1**

Classification performance: CNNs vs. SwinT. Top-1, Top-3 accuracies, weighted  $F_1$  score and MCC are used to evaluate the classification performance of models on the test set of CS data. Bold numbers indicate best performance and increments, while underlined numbers are best baselines for comparisons.

Model	# Para. (M)	Top-1 Acc. (%)	Top-3 Acc. (%)	$F_1$ (weighted)	MCC
WRN-50-2	66.88	65.59	87.63	0.640	0.608
WRN-101-2	124.88	65.52	86.98	<u>0.644</u>	0.608
EfficientNetV2-M	52.88	63.38	85.77	<u>0.623</u>	0.585
EfficientNetV2-L	117.26	65.95	<u>88.34</u>	0.642	0.613
ResNeXt-50	23.88	<u>66.17</u>	86.91	0.642	<u>0.616</u>
ResNeXt-101	81.45	64.45	88.27	0.623	0.596
SwinT-S	48.85	67.53	90.27	0.658	0.632
SwinT-B	86.90	<b>70.03 (+3.86)</b>	90.70	<b>0.690 (+0.046)</b>	<b>0.659 (+0.043)</b>
SwinT-L	195.23	69.74	<b>91.42 (+3.08)</b>	0.688	0.656

The learning rate, weight decay, and the batch size were tuned to  $5 \times 10^{-6}$ , 0.05, and 16 respectively. With respect to the data preprocessing pipeline during training, images were first resized to  $384 \times 384$  and then applied with random crop and rotation transformation before being fed to the model. The model was trained for 50 epochs as the learning performance was saturated by this point. In terms of the hardware, all of our experiments were run on a single NVIDIA 4090 GPU with 24 GB memory.

For hyperparameters unique to SupCon, we followed its original settings proposed in (Khosla et al., 2020). During the pretraining stage, the projection head used was a two-layer MLP with the ReLU activation. The batch size was doubled by independently applying data transformations of random crop and rotation for two times. The temperature in the contrastive loss was set to 0.1. Contrastive pretraining was conducted for 100 epochs. However, we found the transfer performance using the pretrained image encoder plateaued after 50 epochs. Accordingly, the image encoder pretrained for 50 epochs was adopted for the transfer learning experiments.

### 3. Results

Based on our model architecture choices for habitat classification presented in Section 2.2, we first report and discuss the performance of CNN baselines and ViTs in this section, examining both classification performance and model interpretability. Subsequently, we compare SupCon with classical supervised learning described in Section 2.3, focusing on classification performance on challenging habitats that are visually similar. Finally, we evaluate our best-performing model against human experts on a subset of the test data, revealing the performance gap and discussing the potential of our approach for real-world deployment.

#### 3.1. Classification performance

Table 1 compares the classification performance of the chosen CNNs and the ViTs on the habitat classification of the test set. SwinT models outperform all CNN baselines by noticeable margins in all four classification metrics. Specifically, SwinT-B improves Top-1 accuracy,  $F_1$  score, and MCC by 3.86%, 0.046, and 0.043, respectively. Top-3 accuracy is increased by 3.08% with the SwinT-L, exceeding 91%. In terms of the performance of CNN baselines, two WRNs, EfficientNetV2-L, and ResNeXt-50 are on par with each other, with less than 0.6% performance difference on Top-1 accuracy which could be attributed to the noise during training. The close performance observed in the CNNs suggests that this type of model could have reached a performance ceiling on the habitat classification task, which is surpassed by SwinT models.

In addition to the overall performance, we further make detailed pairwise comparison between representatives of the two model architectures on individual habitats. Particularly, ResNeXt-50 and SwinT-B are employed as the representatives. ResNeXt-50 is a strong CNN baseline that achieves the best performance on two out of four classification

metrics, whereas SwinT-B outperforms SwinT-S and SwinT-L on three classification metrics. Table 2 compares SwinT-B with ResNeXt-50 by breaking down the weighted  $F_1$  score shown in Table 1 into precision, recall, and  $F_1$  score on major habitat types. To simplify the results, we grouped L3 performance into L2 categories. For instance, Improved Grassland and Neutral Grassland both belong to the grassland group and each prediction from them contributes to the precision, recall, and the resulting  $F_1$  score of the grassland group. In summary, Table 2 reveals that the SwinT-B surpasses ResNeXt-50 across all individual habitats on  $F_1$  score, ranging between 0.004 to 0.112.

#### 3.2. Prediction interpretability

The merits of employing SwinT models on our habitat classification task are apparent from the results shown in Tables 1 and 2. Furthermore, SwinT exhibits superiority on the interpretability over CNNs by paying attention to broader visual features than its CNN counterparts on habitat images.

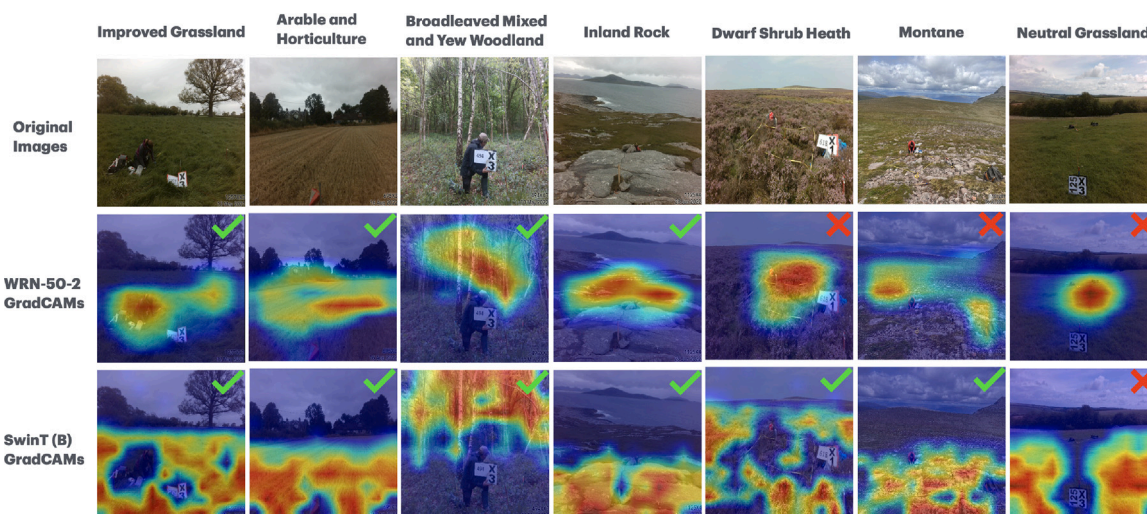
Specifically, we demonstrate the model interpretability with Grad-CAMs. Fig. 5 displays exemplary GradCAMs produced by SwinT-B and WRN-50-2 on different habitats. The observations are twofold: (1) Compared to WRN-50-2, the prediction of SwinT-B is based on a much larger visual context of the images. In contrast, the visual features leveraged by the WRN-50-2 for inference is rather narrow-scoped. The classification made by SwinT-B is much more intuitive as its broad focus is more suitable for making judgement for habitat images. (2) SwinT-B largely ignores distracting visual features, such as on-field surveyors and their equipment, which are irrelevant for habitat classification, whereas, WRN-50-2 is occasionally distracted by such visual features (see the first column and the last column in Fig. 5). Similar GradCAMs to WRN-50-2 are found on other employed CNN models in this study, indicating CNNs could correlate non-determinist visual objects to its predictions, subsequently leading to poor interpretability.

We attribute the illustrated interpretability of SwinT-B for habitat classification to its attention mechanism. Take the grassland as an example, SwinT learns to connect small grass regions shown across the entire image as they are visually similar to each other, and leverage the learned embeddings of connected grass regions to activate its classification layer for the prediction. Conversely, the prediction of CNN is based on confined local areas, without a learning mechanism to connect them over the entire image, resulting in its subjectiveness to local visual features, which makes CNN a poor fit for habitat classification. Combining the observations in classification performance and GradCAMs, we conclude that the capability of ViT for learning image representation encoding long-range regional dependencies and the overall contextual information of the scene, is crucial for improving the model for habitat recognition, and results in superior classification performance over CNN models. While SwinT generalises over a broad scene context, this does not necessarily lead to correct predictions. As shown in Fig. 5 (rightmost column), SwinT misclassifies a Neutral Grassland photo to Improved Grassland. The error arises from the difficulty of distinguishing habitats with highly similar visual context.

**Table 2**

Comparing SwinT-B vs. ResNeXt-50 on individual habitats. This table shows the classification performance when L3 habitats are aggregated into corresponding L2 habitats (to simplify the results presented). Precision and recall are shown along with the  $F_1$  score. Bracketed values represent the performance change of SwinT-B relative to ResNeXt-50.

Coarse class	#Test samples	Model	Precision	Recall	$F_1$
Grasslands	579	ResNeXt-50	<b>0.601</b>	0.603	0.602
		SwinT-B	0.585 (-0.018)	<b>0.629 (+0.027)</b>	<b>0.606 (+0.004)</b>
Croplands	240	ResNeXt-50	0.807	<b>0.942</b>	0.869
		SwinT-B	<b>0.875 (+0.068)</b>	0.929 (-0.013)	<b>0.901 (+0.032)</b>
Wetlands	226	ResNeXt-50	0.590	<b>0.637</b>	0.613
		SwinT-B	<b>0.782 (+0.192)</b>	0.620 (-0.017)	<b>0.691 (+0.045)</b>
Woodlands	200	ResNeXt-50	0.789	0.655	0.716
		SwinT-B	<b>0.801 (+0.012)</b>	<b>0.745 (+0.090)</b>	<b>0.772 (+0.056)</b>
Heathlands	123	ResNeXt-50	0.580	0.561	0.570
		SwinT-B	<b>0.645 (+0.065)</b>	<b>0.724 (+0.163)</b>	<b>0.682 (+0.112)</b>



**Fig. 5.** Comparison of the GradCAMs generated by WRN-50-2 and SwinT-B. The green notation suggests a correct prediction on the sample and the red notation indicates a misclassification. Whether correct or not, the SwinT-B judges the habitat based on a broader and more connected background across the entire image, with a particular overlook on irrelevant visual features, resulting in an improved interpretability over WRN-50-2 on habitat classification task.

Neutral and Improved Grasslands are often subtly different and difficult to interpret in the field. Hence, addressing confusions caused by similar visual context is essential to further improve overall accuracy.

### 3.3. Comparing SupCon with supervised learning

This section compares SupCon with classical supervised learning for training the model with the best reported performance in the previous section, the SwinT-B, and aims to validate the hypothesis that SupCon can improve the classification performance on visually similar habitats.

Table 1 has reported the classification performance achieved by SwinT-B under the supervised learning paradigm. Fig. 6 further shows the classification performance on individual L3 habitats achieved by the supervised learning by illustrating the normalised confusion matrix (CM) produced by SwinT-B. For major habitats, it is observed that the classification performance is relatively modest on grasslands, such as Neutral Grassland, Acid Grassland, and Bracken, as well as on Fen, Marsh, Swamp (FMS) belonging to the wetland. Less than 70% of test instances among these habitats are correctly predicted by SwinT-B (shown by the diagonal values in the corresponding rows of CM). This observation indicates that these habitats are very challenging to classify even using ViT (CNNs perform closely to ViT or even worse on these habitats in our experiments). We attribute this performance degradation to the visual similarity between certain habitats, as illustrated in Fig. 1. For instance, Neutral Grassland, which is a major L3 habitat in the CS dataset and often displays subtle visual differences to Improved Grassland, has 34% of its test instances being misclassified as Improved

Grassland, deteriorating the overall performance significantly. In Appendix A, we visualise the embedding clusters of different habitats. The heavily overlapped clusters further reflect their high similarity.

We applied SupCon to the top-performing variant from each CNN type (WRN-50-2, EfficientNetV2-L, ResNeXt-50) and to the strongest Swin Transformer (SwinT-B), as identified in our supervised learning comparison, using identical contrastive pretraining settings. Table 3 reports their classification performance, denoted with the changes based on the classical supervised learning indicated in Table 1. SupCon slightly improves the performance of SwinT-B on all four metrics (+1.50% Top-1, +0.22% Top-3, +0.016  $F_1$ , +0.018 MCC). Regarding CNNs, SupCon gives marginal gains for EfficientNetV2-L across metrics and has mixed effects on WRN-50-2 and ResNeXt-50 with slight improvements in Top-3 but degradations in Top-1/ $F_1$ /MCC. These results suggest that, SupCon provides the clearest benefit for the Swin Transformer under the same settings, with inconsistent performance for the CNN baselines.

For our best model, SwinT-B, SupCon is observed to mitigate the misclassifications between challenging habitats. Specifically, Fig. 7 is a delta confusion matrix that compares the normalised confusion matrices generated by SupCon and supervised learning on major habitats. With SupCon, Improved Grassland increases its correct classifications by 6% and reduces the misclassifications on Neutral Grassland by 6%. Similarly, 5% more Neutral Grassland instances are correctly predicted and its misclassifications to Improved Grassland are reduced by 2% with SupCon. In terms of FMS, SupCon improves the proportion of correct classifications by 14% and decrease its misclassifications to



**Fig. 6.** The confusion matrix of SwinT-B on the test set. Habitats with no samples in the test set are removed from the matrix (refer to Section 2.5 for details). While SwinT-B demonstrates more suitability than CNNs for habitat classification, some habitats are often misclassified by the model, e.g., Neutral Grassland is often misclassified as Improved Grassland, affecting the overall classification performance significantly.

**Table 3**

Performance of supervised learning (Sup) vs. supervised contrastive learning (SupCon) across architectures. Numbers are reported on the CS habitat test set. For each model, SupCon cells show the absolute gain over Sup in parentheses.

Model	Top-1 Acc. (%)	Top-3 Acc. (%)	$F_1$ (weighted)	MCC
WRN-50-2	64.23 (-1.36) ↓	89.13 (+1.50) ↑	0.627 (-0.013) ↓	0.593 (-0.015) ↓
EfficientNetV2-L	67.02 (+1.07) ↑	90.70 (+2.36) ↑	0.658 (+0.016) ↑	0.625 (+0.012) ↑
ResNeXt-50	64.23 (-1.94) ↓	88.20 (+1.29) ↑	0.610 (-0.032) ↓	0.593 (-0.023) ↓
SwinT-B	<b>71.53</b> (+1.50) ↑	<b>90.92</b> (+0.22) ↑	<b>0.706</b> (+0.016) ↑	<b>0.677</b> (+0.018) ↑

Bog and Neutral Grassland by 5% and 6%, respectively. The correct predictions of Bracken are improved by 8% with SupCon, owing to a 12% decrease of misclassifications to Acid Grassland. Conversely, SupCon sacrifices classification performance of Dwarf Shrub Heath, decreasing its correct predictions by 4%, and fails to increase the correct classifications for Calcareous Grassland. This could be due to the fact that the Calcareous Grassland is a minor L3 habitat with a very limited number of samples in CS dataset for training the model, as shown in Fig. 2. We provide more insights about the performance of SupCon with the embedding analysis in Appendix A.

For our CNN backbones, we observe that SupCon can shift the decision boundary for them rather than enlarge the inter-class margin: for Neutral vs Improved Grassland, SupCon reduces Neutral to Improved confusions but increases Improved to Neutral, yielding higher recall for Neutral at the expense of precision. A plausible explanation is the

strong local prior of CNNs, which may limit the visual information expressed in the embeddings, making contrastive learning less effective to separate classes that require broader contextual cues. In contrast, SwinT-B with SupCon reduces errors in both directions, suggesting that global contextual cues leveraged by the SwinT-B are better suited to contrastive learning for the habitat classification task. Further details on the performance of CNNs with SupCon are provided in the Appendix B.

### 3.4. Comparing the model with human experts

Table 4 compares the performance of our best performed CNN and ViT trained with SupCon as shown in Table 3, on a subset of the test data with photographic interpretation of L3 habitats by three human experts, as specified in Section 2.4.3. The notably higher performance

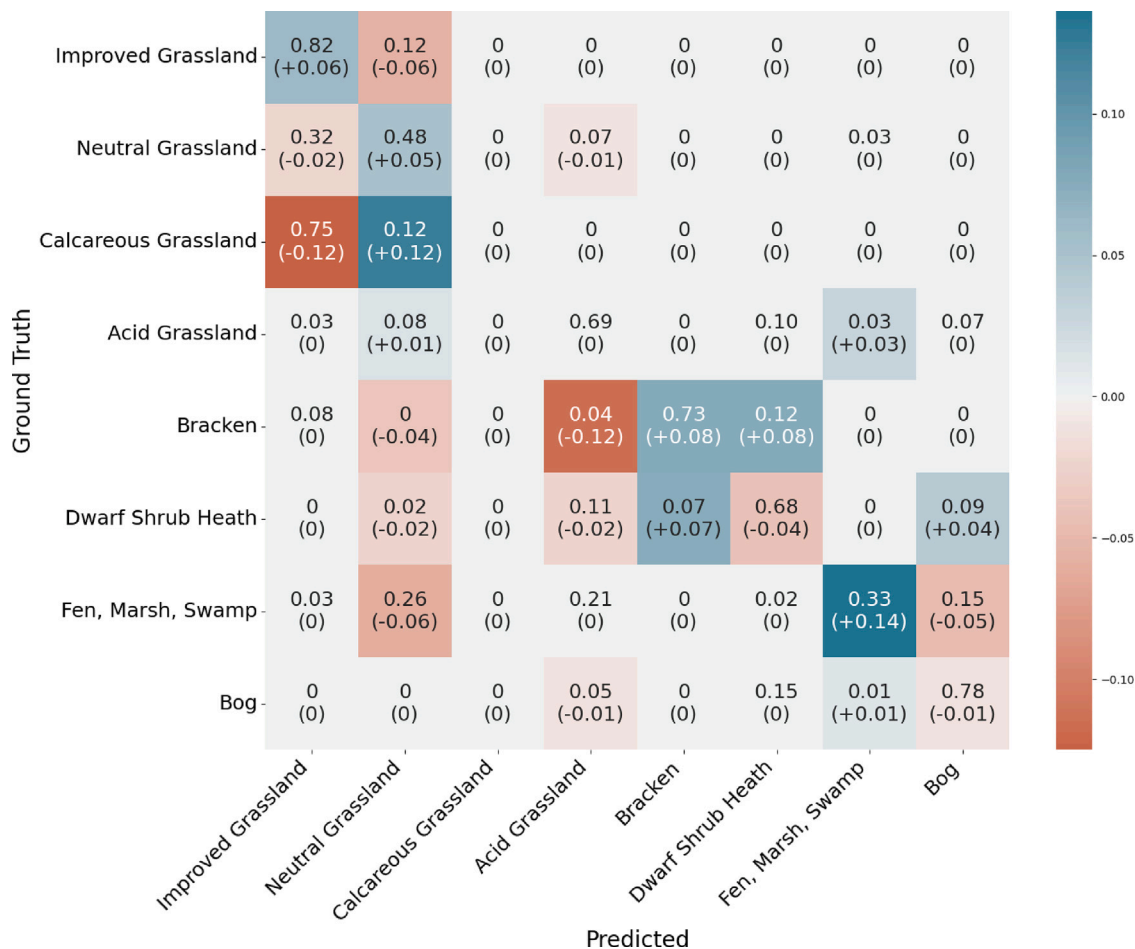


Fig. 7. Delta confusion matrix (CM) of challenging habitats: based on the CM produced by SupCon, this graph further highlights its differences from the CM generated by supervised learning, with blue colour indicating boost and red colour indicating decrease. SupCon reduces misclassifications on three major grasslands and the FMS in wetlands, which confuse the model most in the supervised learning.

Table 4

Performance comparison between our best models, EfficientNetV2-L and SwinT-B, and human experts on a subset drawn from the test set. Bold values are the best performance, underlined values are the best performance achieved by human experts, and performance differences between the model and the best human performance are indicated in the brackets.

Metric	EfficientNetV2-L (SupCon)	SwinT-B (SupCon)	Expert 1	Expert 2	Expert 3
Top-1 Acc. (%)	56.33 (-7.36)	<b>63.92 (+0.23)</b>	55.70	51.27	<u>63.69</u>
Top-3 Acc. (%)	84.18 (+4.56)	<b>87.34 (+7.72)</b>	73.42	74.68	<u>79.62</u>
$F_1$ (weighted)	0.535 (-0.093)	0.603 (-0.025)	0.541	0.529	<u>0.628</u>
MCC vs Ground Truth	0.516 (-0.088)	0.599 (-0.005)	0.517	0.479	<u>0.604</u>
MCC vs Expert 1	0.529	0.529	-	-	-
MCC vs Expert 2	0.460	0.490	0.515	-	-
MCC vs Expert 3	0.530	0.546	0.511	0.553	-

of Expert 3 likely reflects their extensive familiarity with the test data, gained through years of participation in the UK Countryside Survey. Compared to the best expert (Expert 3), SwinT-B: (1) outperforms their Top-1 accuracy by 0.23% and their Top-3 accuracy significantly by 7.72%; (2) slightly underperforms on weighted  $F_1$  and MCC by 0.025 and 0.005, respectively, primarily because it struggles to generalise to

rare habitats (e.g. Montane) owing to limited training samples (see Appendix C for detailed discussion); (3) has a strong agreement with the best expert, achieving an MCC of 0.546 against them. In contrast, our best CNN with SupCon (EfficientNetV2-L) underperforms the human experts on all headline metrics and is well below SwinT-B, indicating that studied CNNs do not yet reach expert-level performance on this task.

#### 4. Discussion

Our experiment results show that SwinT surpasses state-of-the-art CNNs, including Wide ResNet, EfficientNetV2, and ResNeXt, for ground-level habitat classification. SupCon training can further boost SwinT, achieving performance comparable to human experts. Meanwhile, the study also exposes challenges in habitat classification with ground-level imagery. We provide the following insights:

- **Inherent challenge of ground-level imagery.** With a Top-1 accuracy of only ~63% in human benchmark, our study reveals that ground-level classification of L3 habitats from images remains difficult for both AI models and human experts. Remote sensing methods have shown better performance in multi-habitat (albeit narrower) contexts, e.g. forests (Praticò et al., 2021), wetlands (Jamali et al., 2023), grasslands (Díaz-Ireland et al., 2024). Unlike bird's-eye remote-sensing data, ground-level images require models to understand scene depth (disentangling background from foreground), to filter out distracting elements present in the scene, and to detect often-subtle compositional differences among visually similar subtypes. These added complexities make ground-level habitat classification very challenging. Misclassification by the AI model can trigger inappropriate actions on the ground. From an ecological perspective, misclassified habitats could lead to wrong interventions or recommendations for land management, e.g. applying mineral fertiliser to neutral/species-rich grassland. For Biodiversity Net Gain, inaccurate mapping could result in development on ecologically valuable habitats or inappropriate decisions about where to offset habitat losses following development. Therefore, addressing the challenges of ground-level image classification is critical to responsible habitat monitoring.
- **Pathways to improving ground-level habitat classification.** (a) The CS photos were originally taken for plot relocation, not habitat identification. Establishing clear guidelines (e.g. framing, lighting, seasonal timing) for habitat-focused imagery should enhance the training data and reveal the true potential of AI for ground-level habitat classification. (b) Human surveyors use a rich collection of information and often quantitative data—information absent from 2D images—to determine habitats. For example, they take into consideration the geographical and climatic information, management context, soil type, species, etc. when conducting habitat assessment in the field. Additionally, they may visit the site during different seasons in order to record species growing at different times of the year. Incorporating multi-view or time-series photos, or auxiliary metadata (e.g. date, location), could narrow the remaining performance gap. (c) Since habitat is often defined by characteristic species, extending our approach to recognise key flora (e.g. dominant shrubs, indicator plants) and integrating those outputs into the habitat classifier may also boost accuracy.
- **Practical deployment.** When deploying the model in practice, a well-designed sampling protocol is critical to ensure reliable performance. We propose: capturing clear, context-rich images with the rear camera supported by automatic quality checks; sampling multiple photo-points within each spatial unit; and acquiring several views per point (forward, nadir, and detail). Predictions are then aggregated to habitat classes with uncertainty estimates and flags for expert review, ensuring responsible use. From a user-experience perspective, network latency is the principal constraint: a reliable connection is required to transmit images and return results with minimal delay. By contrast, compute time is comparatively negligible: on our workstation (NVIDIA RTX 4090), a single SwinT-B forward pass takes milliseconds. In our web-based API prototype, end-to-end latency is typically 0.5–2.0 s per image.

#### 5. Conclusion

In this work, we investigated deep learning approaches for habitat classification based on ground-level imagery to bridge the research gaps in AI-driven habitat monitoring. Guided by the unique challenges of ground habitat scenes (broad and continuous regions, visually similar habitats), we compared two deep neural architectures, CNNs and ViTs, and two training paradigms, standard supervised learning vs supervised contrastive learning (SupCon).

Our evaluation shows that ViTs outperform state-of-the-art CNN baselines—achieving up to a 4% increase in overall accuracy and similar gains across other key metrics. Beyond raw accuracy, ViTs generalise with a broader habitat scene understanding, demonstrating superior interpretability over CNNs. We also find that SupCon is a more effective training paradigm than classical supervised learning for training ViT to distinguish between visually similar habitats. SupCon considerably reduces misclassifications between commonly confused habitat pairs—such as Improved and Neutral Grassland—outperforming standard supervised learning. Finally, our best ViT model not only exceeds CNN performance but also matches photographic identification by experienced ecological experts on key classification metrics, highlighting the potential of our work to reduce the need for human experts to interpret habitats at this level in the field.

Our current study focuses solely on level 3 Broad Habitats within the UK (except for Bracken, which is L4). Future work should expand the model to classify fine-grained level 4 habitats with over 80 categories, where sample sizes are still small, and assess model transferability to other countries and different habitats outside of the UK. In summary, our study paves a path towards truly scalable, cost-effective habitat monitoring based on ground-level imagery by integrating state-of-the-art AI techniques with ecological expertise to support better-informed biodiversity conservation and land-use decisions.

#### CRedit authorship contribution statement

**Hongrui Shi:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Lisa Norton:** Writing – review & editing, Validation, Supervision, Resources, Project administration, Funding acquisition, Data curation, Conceptualization. **Lucy Ridding:** Writing – review & editing, Validation, Resources, Project administration, Data curation. **Simon Rolph:** Writing – review & editing, Resources, Project administration. **Tom August:** Writing – review & editing, Resources, Funding acquisition. **Claire M. Wood:** Resources, Data curation. **Lan Qie:** Writing – review & editing, Validation, Supervision, Project administration, Funding acquisition, Conceptualization. **Petra Bosilj:** Writing – review & editing, Supervision, Resources, Project administration, Methodology, Funding acquisition, Conceptualization. **James M. Brown:** Writing – review & editing, Supervision, Resources, Project administration, Methodology, Funding acquisition, Conceptualization.

#### Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work the author(s) used ChatGPT in order to improve language and readability. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the content of the publication.

#### Funding

This work was funded by Natural Environment Research Council of UK Research and Innovation [grant numbers UKRI054].

## Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.ecoinf.2026.103751>.

## Data availability

The code used for model training and evaluation is publicly available at <https://github.com/WhiteGiveFive/Habitat-Classification-from-Ground-Level-Imagery> (MIT license). The data used in this study are not publicly hosted due to UKCEH's data privacy policy but may be obtained from UKCEH under a research licence upon request. Requests for access to the ground-level photographs and corresponding UKHab labels should be directed to Dr Lisa Norton ([lrn@ceh.ac.uk](mailto:lrn@ceh.ac.uk)).

## References

- Bjerge, K., Mann, H.M., Høye, T.T., 2022. Real-time insect tracking and monitoring with computer vision and deep learning. *Remote. Sens. Ecol. Conserv.* 8 (3), 315–327.
- Breiman, L., 2001. Random forests. *Mach. Learn.* 45, 5–32.
- Carey, P.D., Wallis, S., Emmett, B., Maskell, L., Murphy, J., Norton, L., Simpson, I., Smart, S., 2008. Countryside Survey: UK Headline Messages from 2007. NERC/Centre for Ecology & Hydrology.
- Cortes, C., Vapnik, V., 1995. Support-vector networks. *Mach. Learn.* 20, 273–297.
- Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L., 2009. Imagenet: A large-scale hierarchical image database. In: 2009 IEEE Conference on Computer Vision and Pattern Recognition. IEEE, pp. 248–255.
- Díaz-Ireland, G., Gülçin, D., López-Sánchez, A., Pla, E., Burton, J., Velázquez, J., 2024. Classification of protected grassland habitats using deep learning architectures on sentinel-2 satellite imagery data. *Int. J. Appl. Earth Obs. Geoinf.* 134, 104221.
- Dosovitskiy, A., Beyler, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al., 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- zu Ermgassen, S.O., Marsh, S., Ryland, K., Church, E., Marsh, R., Bull, J.W., 2021. Exploring the ecological outcomes of mandatory biodiversity net gain using evidence from early-adopter jurisdictions in England. *Conserv. Lett.* 14 (6), e12820.
- European Commission, 2020. EU Biodiversity Strategy for 2030: Bringing Nature Back into Our Lives. Commission Communication COM(2020) 380 final, URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex:52020DC0380>.
- Gómez-Ríos, A., Tabik, S., Luengo, J., Shihavuddin, A., Krawczyk, B., Herrera, F., 2019. Towards highly accurate coral texture images classification using deep convolutional neural networks and data augmentation. *Expert Syst. Appl.* 118, 315–328.
- Goodfellow, I., Bengio, Y., Courville, A., Bengio, Y., 2016. *Deep Learning*, vol. 1, no. 2, MIT press Cambridge.
- GOV.UK, 2023. Biodiversity Net Gain. Department for Environment, Food & Rural Affairs, URL: <https://www.gov.uk/government/collections/biodiversity-net-gain..>
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 770–778.
- Jamali, A., Roy, S.K., Ghamisi, P., 2023. WetMapFormer: A unified deep CNN and vision transformer for complex wetland mapping. *Int. J. Appl. Earth Obs. Geoinf.* 120, 103333.
- Joint Nature Conservation Committee, 2025. UK BAP priority habitats. URL: <https://jncc.gov.uk/our-work/uk-bap-priority-habitats/>. (Accessed 01 July 2025).
- Joly, A., Picek, L., Kahl, S., Goëau, H., Espitalier, V., Botella, C., Marcos, D., Estopinan, J., Leblanc, C., Larcher, T., et al., 2024. Overview of lifeclif 2024: Challenges on species distribution prediction and identification. In: International Conference of the Cross-Language Evaluation Forum for European Languages. Springer, pp. 183–207.
- Khosla, P., Teterwak, P., Wang, C., Sarna, A., Tian, Y., Isola, P., Maschinot, A., Liu, C., Krishnan, D., 2020. Supervised contrastive learning. *Adv. Neural Inf. Process. Syst.* 33, 18661–18673.
- Krizhevsky, A., Hinton, G., et al., 2009. Learning multiple layers of features from tiny images.
- Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B., 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 10012–10022.
- Liu, Z., Mao, H., Wu, C.Y., Feichtenhofer, C., Darrell, T., Xie, S., 2022. A convnet for the 2020s. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11976–11986.
- Loshchilov, I., Hutter, F., 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Mall, U., Phoo, C.P., Liu, M.K., Vondrick, C., Hariharan, B., Bala, K., 2023. Remote sensing vision-language foundation models without annotations via ground remote alignment. *arXiv:2312.06960*.
- Marcinkowska-Ochtyra, A., Ochtyra, A., Raczek, E., Kopeć, D., 2023. Natura 2000 grassland habitats mapping based on spectro-temporal dimension of sentinel-2 images with machine learning. *Remote. Sens.* 15 (5), 1388.
- Marjani, M., Mohammadimanesh, F., Mahdianpari, M., Gill, E.W., 2025. A novel spatio-temporal vision transformer model for improving wetland mapping using multi-seasonal sentinel data. *Remote. Sens. Appl.: Soc. Environ.* 37, 101401.
- Martínez-Sánchez, L., See, L., Yordanov, M., Verhegghen, A., Elvekjær, N., Muraro, D., d'Andrimont, R., Van der Velde, M., 2024. Automatic classification of land cover from LUCAS in-situ landscape photos using semantic segmentation and a random forest model. *Environ. Model. Softw.* 172, 105931.
- Morueta-Holme, N., Iversen, L., Corcoran, D., Rahbek, C., Normand, S., 2024. Unlocking ground-based imagery for habitat mapping. *Trends Ecol. Evolut.* 39 (4), 349–358.
- Norris, G.S., LaRocque, A., Leblon, B., Barbeau, M.A., Hanson, A.R., 2024. Comparing pixel- and object-based approaches for classifying multispectral drone imagery of a salt marsh restoration and reference site. *Remote. Sens.* 16 (6), 1049.
- Oba, Y., Doi, H., 2025. Accelerating ecosystem monitoring through computer vision with deep metric learning. *Ecol. Complex.* 62, 101124.
- Perrett, A., Pollard, H., Barnes, C., Schofield, M., Qie, L., Bosilj, P., Brown, J.M., 2023. DeepVerge: classification of roadside verge biodiversity and conservation potential. *Comput. Environ. Urban Syst.* 102, 101968.
- Pl@ntNet Consortium, 2025. Pl@ntNet: Plant identification platform. <https://plantnet.org/>. (Accessed 14 May 2025).
- Praticò, S., Solano, F., Di Fazio, S., Modica, G., 2021. Machine learning classification of Mediterranean forest habitats in google earth engine based on seasonal sentinel-2 time-series and input image composition optimisation. *Remote. Sens.* 13 (4), 586.
- Ratnayake, M.N., Dyer, A.G., Dorin, A., 2021. Tracking individual honeybees among wildflower clusters with computer vision-facilitated pollinator monitoring. *PLoS One* 16 (2), e0239504.
- Retallack, A., Finlayson, G., Ostendorf, B., Lewis, M., 2022. Using deep learning to detect an indicator arid shrub in ultra-high-resolution UAV imagery. *Ecol. Indic.* 145, 109698.
- Reynolds, S.A., Beery, S., Burgess, N., Burgman, M., Butchart, S.H., Cooke, S.J., Coomes, D., Danielsen, F., Di Minin, E., Durán, A.P., et al., 2025. The potential for AI to revolutionize conservation: a horizon scan. *Trends Ecol. Evolut.* 40 (2), 191–207.
- Saadeldin, M., O'Hara, R., Zimmermann, J., Mac Namee, B., Green, S., 2022. Using deep learning to classify grassland management intensity in ground-level photographs for more automated production of satellite land use maps. *Remote. Sens. Appl.: Soc. Environ.* 26, 100741.
- Schofield, L.R., Pearson, M.E., Newell, S., Clackum, N., Turner, B.L., 2024. Why aren't more landowners enrolling in land-based carbon credit exchanges? *Rangelands* 46 (4), 117–131.
- Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D., 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 618–626.
- Simonyan, K., Zisserman, A., 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Sittaro, F., Hutengs, C., Semella, S., Vohland, M., 2022. A machine learning framework for the classification of natura 2000 habitat types at large spatial scales using MODIS surface reflectance data. *Remote. Sens.* 14 (4), 823.
- Tan, M., Le, Q., 2021. Efficientnetv2: Smaller models and faster training. In: International Conference on Machine Learning. PMLR, pp. 10096–10106.
- U.K. Biodiversity Group, 2008. UK biodiversity action plan priority habitat descriptions. URL: <https://hub.jncc.gov.uk/assets/0b7943ea-2eeec-47a9-bd13-76d1d66d471f..>
- UKHab Ltd, 2023. UK habitat classification system. URL: <https://www.ukhab.org/Professional and Basic editions; registration required for full documentation>.
- Van An, N., Quang, N.H., Son, T.P.H., An, T.T., 2023. High-resolution benthic habitat mapping from machine learning on PlanetScope imagery and ICESat-2 data. *Geocarto Int.* 38 (1), 2184875.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I., 2017. Attention is all you need. *Adv. Neural Inf. Process. Syst.* 30.
- Wang, Q., 2025. Plant species recognition and classification algorithm based on deep learning. In: 2025 3rd International Conference on Integrated Circuits and Communication Systems. ICICACS, IEEE, pp. 1–5.
- Wang, W., Dai, J., Chen, Z., Huang, Z., Li, Z., Zhu, X., Hu, X., Lu, T., Lu, L., Li, H., et al., 2023. Internimage: Exploring large-scale vision foundation models with deformable convolutions. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14408–14419.
- Wood, C.M., Smart, S.M., Bunce, R.G., Norton, L.R., Maskell, L.C., Howard, D.C., Scott, W.A., Henrys, P.A., 2017. Long-term vegetation monitoring in Great Britain—the countryside survey 1978–2007 and beyond. *Earth Syst. Sci. Data* 9 (2), 445–459.

- Wu, Z., Xiong, Y., Yu, S.X., Lin, D., 2018. Unsupervised feature learning via non-parametric instance discrimination. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3733–3742.
- Xie, S., Girshick, R., Dollár, P., Tu, Z., He, K., 2017. Aggregated residual transformations for deep neural networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1492–1500.
- Xu, G., Xie, Y., Luo, Y., Yin, Y., Li, Z., Wei, Z., 2024. Advancing automated surveillance: Real-time detection of crown-of-thorns starfish via YOLOv5 deep learning. *J. Theory Pr. Eng. Sci.* 4 (06), 1–10.
- Zagoruyko, S., Komodakis, N., 2016. Wide residual networks. arXiv preprint arXiv:1605.07146.
- Zhong, J., Li, M., Zhang, H., Qin, J., 2023. Fine-grained 3D modeling and semantic mapping of coral reefs using photogrammetric computer vision and machine learning. *Sensors* 23 (15), 6753.
- Zhu, L., Wang, X., Ke, Z., Zhang, W., Lau, R.W., 2023. Biformer: Vision transformer with bi-level routing attention. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10323–10333.