

Article

AI-Driven Real-Time Monitoring of Ground-Nesting Birds: A Case Study on Curlew Detection Using YOLOv10

Carl Chalmers ^{1,*}, Paul Fergus ¹, Serge Wich ¹, Steven N. Longmore ¹, Naomi Davies Walsh ¹, Lee Oliver ², James Warrington ², Julieanne Quinlan ² and Katie Appleby ²

¹ Faculty of Health, Innovation, Technology and Science, Liverpool John Moores University, Liverpool L3 3AF, UK; p.fergus@ljmu.ac.uk (P.F.); s.wich@ljmu.ac.uk (S.W.); s.n.longmore@ljmu.ac.uk (S.N.L.); n.j.walsh@2022.ljmu.ac.uk (N.D.W.)

² Game & Wildlife Conservation Trust, Fordingbridge SP6 1EF, UK; loliver@gwct.org.uk (L.O.); jwarrington@gwct.org.uk (J.W.); jquinlan@gwct.org.uk (J.Q.); kappleby@gwct.org.uk (K.A.)

* Correspondence: c.chalmers@ljmu.ac.uk

Abstract: Effective monitoring of wildlife is critical for assessing biodiversity and ecosystem health as declines in key species often signal significant environmental changes. Birds, particularly ground-nesting species, serve as important ecological indicators due to their sensitivity to environmental pressures. Camera traps have become indispensable tools for monitoring nesting bird populations, enabling data collection across diverse habitats. However, the manual processing and analysis of such data are resource-intensive, often delaying the delivery of actionable conservation insights. This study presents an AI-driven approach for real-time species detection, focusing on the curlew (*Numenius arquata*), a ground-nesting bird experiencing significant population declines. A custom-trained YOLOv10 model was developed to detect and classify curlews and their chicks using 3/4G-enabled cameras linked to the Conservation AI platform. The system processes camera trap data in real time, significantly enhancing monitoring efficiency. Across 11 nesting sites in Wales, the model achieved high performance, with a sensitivity of 90.56%, specificity of 100%, and F1-score of 95.05% for curlew detections and a sensitivity of 92.35%, specificity of 100%, and F1-score of 96.03% for curlew chick detections. These results demonstrate the capability of AI-driven monitoring systems to deliver accurate, timely data for biodiversity assessments, facilitating early conservation interventions and advancing the use of technology in ecological research.

Keywords: conservation; object detection; image processing; modelling biodiversity; deep learning; camera traps; *Numenius arquata*



Academic Editors: Margarita Mulero-Pazmany and Maria Laura Carranza

Received: 21 November 2024

Revised: 13 February 2025

Accepted: 21 February 2025

Published: 23 February 2025

Citation: Chalmers, C.; Fergus, P.; Wich, S.; Longmore, S.N.; Walsh, N.D.; Oliver, L.; Warrington, J.; Quinlan, J.; Appleby, K. AI-Driven Real-Time Monitoring of Ground-Nesting Birds: A Case Study on Curlew Detection Using YOLOv10. *Remote Sens.* **2025**, *17*, 769. <https://doi.org/10.3390/rs17050769>

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Ground-nesting birds are among the most vulnerable wildlife groups, facing an array of environmental pressures that threaten their survival [1]. These species, which include lapwings, skylarks, and curlews, rely on open landscapes to breed, where their nests and chicks are particularly exposed to threats such as habitat loss and predation [2]. Agricultural intensification, urban development, and changes in land use have significantly reduced the availability of suitable nesting sites, leaving many ground-nesting birds struggling to maintain stable populations [3]. Predation by meso-predators such as foxes, badgers, and corvids further exacerbate the challenge, often leading to high rates of egg and chick mortality [4]. Ground-nesting birds play an important ecological role in the maintenance of grassland and wetland ecosystems through the regulation of insect populations and

the construction of ground nests that encourage diverse vegetation structures [5,6]. Their nesting activities contribute to soil aeration and nutrient redistribution, fostering plant diversity and supporting a range of other wildlife species. The conservation of these ground-nesting birds is essential not only for their survival but also for the overall health and resilience of the ecosystems they inhabit.

Among these species, the curlew (*Numenius arquata*) is of particular concern, exemplifying the challenges faced by ground-nesting birds [7]. Once widespread across the UK, curlew populations have experienced dramatic declines in recent decades, driven by habitat degradation and predation [8]. This has left the species on the brink of regional extinction, with urgent conservation efforts needed to reverse this trend [9].

Efforts to protect vulnerable bird populations have largely focused on habitat restoration, predator control, and traditional nest monitoring [10]. While these methods have proven valuable, they face significant limitations that hinder their effectiveness. Traditional nest monitoring, which typically involves physical observation or manual review of historical camera trap data, is labour-intensive, time-consuming, and logistically challenging—particularly for species like curlews that nest in remote or hard-to-access areas [11].

The use of camera traps provides a non-invasive and highly effective method for wildlife monitoring, allowing researchers to collect extensive datasets over prolonged periods and across diverse geographical locations [12]. These devices are instrumental in studying animal behaviour, population dynamics, and habitat utilisation. Camera trap usage is particularly important for monitoring ground-nesting birds, which face heightened vulnerability. Camera traps facilitate continuous observation of nests without causing disturbance, offering critical insights into breeding success, nesting behaviours, and interactions with predators. However, the process of manually reviewing and analysing the vast volume of images generated remains a significant limitation [13]. This labour-intensive task delays the implementation of timely conservation measures, thereby hindering the ability to respond effectively to threats [14].

While AI-driven solutions exist to automatically analyse camera trap data, many are designed to filter out blank images and broadly categorise the presence or absence of animals, people, or vehicles, rather than focus on detailed species-level classifications [15]. Models used for species classification often prioritise handling a large number of classes across diverse taxa rather than tailoring their application to specific geographical regions or ecosystems. This approach is particularly problematic for biodiversity monitoring, where distinguishing between visually similar species is essential for targeted conservation efforts [16]. As a result, existing solutions frequently overlook common ground-nesting bird species, making them inadequate for effectively monitoring these vulnerable populations.

Furthermore, many solutions lack real-time data processing. Camera trap images are often stored on SD cards and retrieved manually, resulting in significant delays between data collection and analysis [17]. This workflow is impractical for species like curlews, where immediate interventions may be necessary to protect eggs or chicks from predation. Such interventions include the removal of eggs and chicks for incubation. Some AI systems, such as MegaDetector [18], PyTorch Wildlife [19], Wildlife insights [20], and INaturalist [21], have demonstrated potential in automating aspects of wildlife monitoring. However, their effectiveness is often constrained by reliance on pre-processed datasets, imbalanced class representation [22], inability to generalise across ecosystems, and a lack of integration with real-time data streams [23,24]. In addition, landmark projects such as Snapshot Serengeti [22] have highlighted the effectiveness of DL models for large-scale species classification. Snapshot Serengeti, for instance, used convolutional neural networks (CNNs) to classify millions of camera trap images, significantly reducing manual labour. Despite these achievements, existing research has revealed challenges common in biodiversity

monitoring, such as imbalanced datasets, poor generalisability across ecosystems, and reduced accuracy for rarer or visually similar species. These limitations, coupled with the lack of real-time integration, underscore the need for AI solutions that can deliver species-specific classifications with high accuracy.

Most existing systems prioritise reducing data volume rather than enabling proactive responses, making them ill-suited for dynamic conservation scenarios where timely decision-making is crucial [25]. To address these challenges, there is a need for AI-enabled solutions that combine high accuracy, species-specific classification, and real-time processing capabilities. By automating species detection and providing immediate insights to researchers, such systems can help overcome the limitations of traditional methods and existing AI approaches. These innovations have the potential to revolutionise the way we monitor and protect ground-nesting birds, therefore, enabling more efficient and effective conservation strategies.

In this study, we present an automated object detection and classification pipeline capable of identifying curlews and curlew chicks (Figure 1). Our model utilises a custom-trained YOLOv10 Deep Learning (DL) model, which is integrated with the Conservation AI platform [26]. This model is capable of detecting and classifying curlews and their chicks received from real-time 3/4G-enabled camera traps. The approach automates species detection and supports large-scale biodiversity surveys to enable more timely and accurate ecological assessments, especially for vulnerable species like the curlew. By focusing on curlew monitoring, this study demonstrates the applicability of a scalable, efficient solution for tracking curlews over extended periods to identify key opportunities for intervention. Automating the analysis of camera trap data in this way accelerates data processing while providing real-time alerts on ecologically significant events, such as nesting activity or chick presence, enabling further analysis for timely interventions to support curlew population recovery.



Figure 1. Acquired image of a curlew at one of the locations in the study.

2. Materials and Methods

2.1. Data Collection and Description

The dataset used for modelling contains images of 26 distinct species and objects found in the UK: person, red fox (*Vulpes vulpes*), European fallow deer (*Dama dama*), roe deer (*Capreolus capreolus*), European hedgehog (*Erinaceus europaeus*), wood grouse (*Capercaillie cock*), wood grouse (*Capercaillie hen*), cattle (*Bos taurus*), domestic dog (*Canis*

familiaris), red deer (*Cervus elaphus*), European rabbit (*Oryctolagus cuniculus*), European badger (*Meles meles*), common buzzard (*Buteo buteo*), northern goshawk (*Accipiter gentilis*), domestic cat (*Felis catus*), eastern grey squirrel (*Sciurus carolinensis*), red squirrel (*Sciurus vulgaris*), European pine martin (*Martes martes*), common pheasant (*Phasianus colchicus*), house sparrow (*Passer domesticus*), domestic sheep (*Ovis aries*), common wood pigeon (*Columba palumbus*), common curlew (*Numenius arquata*), common curlew chick (*Numenius arquata*), domestic goat (*Capra hircus*) and calibration pole, which were obtained through various conservation partners and private camera deployments. The dataset used in this study comprised a total of 38,740 image files. The average image resolution across the dataset was 972×769 pixels. This value reflects the characteristics of the training dataset used in this study. An analysis of the resolution distribution revealed no significant outliers that would adversely affect model training as shown in Figure 2. As a result, no images were excluded prior to tagging. Conducting this input resolution analysis was essential for determining the appropriate aspect ratio coefficient, which was then incorporated into the hyperparameter configuration for training the model.




Figure 2. Analysis of image resolution distribution within the dataset identifying outliers that could potentially affect model training and performance.

2.2. Data Pre-Processing

The data annotation process utilised the Conservation AI tagging platform, where bounding boxes were created to define regions of interest within each image. These annotations were exported in XML format following the Pascal VOC standard. Images identified as low quality or unsuitable for model training were categorised as “no good” and excluded from the dataset. The object counts per class ranged from approximately 1100 to 2500, leading to class imbalance (Figure 3) due to uneven representation in the original dataset. Overall, a total of 38,740 objects were annotated across the dataset.

The labelled data were converted to the YOLO annotation format using a Python 3.8 script. The dataset, consisting of images and corresponding labels, was randomly split into 80% for training, 10% for validation, and 10% for testing, based on the tagged annotations. The XML files were converted directly into YOLO-specific text files, where each object is represented by its class and bounding box coordinates.

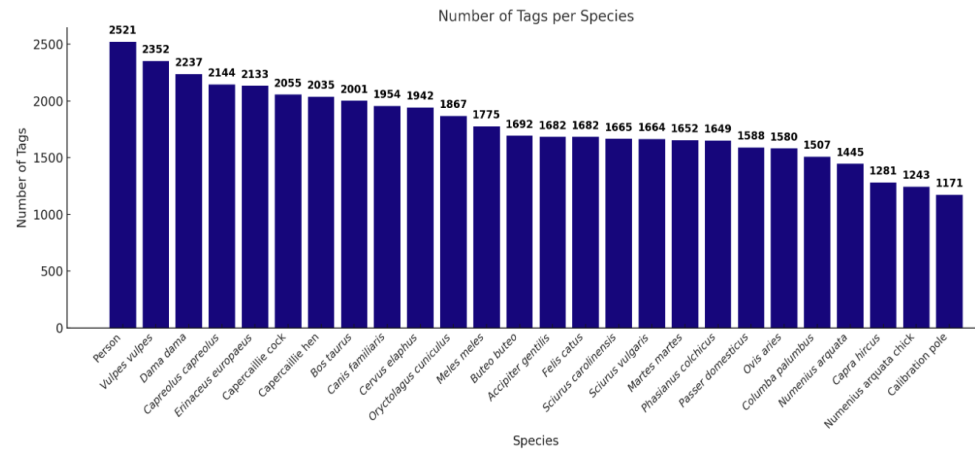


Figure 3. Class distribution of annotated data, with species names represented on the x-axis and the number of tags (bounding boxes) per species shown on the y-axis.

2.3. Model Selection

The YOLOv10x model was employed for both object detection and classification across 26 different species and objects. The model features 29.5 million parameters, balancing high model capacity with computational efficiency, which allows it to perform complex detection tasks at scale without excessive hardware requirements. Unlike the two-stage approach used in models such as Faster-RCNN, YOLOv10x integrates detection and classification in a single-stage process, significantly improving both speed and efficiency. This model utilises the CSPDarknet backbone, which extracts essential feature maps, and incorporates a Path Aggregation Network (PAN) for feature fusion, improving the detection of objects at multiple scales. YOLOv10x further refines object detection with its anchor-free detection mechanism, eliminating the need for predefined anchor boxes, and leverages dynamic convolution to enhance classification accuracy [27]. By predicting bounding boxes and class probabilities simultaneously, the model achieves high-speed inference while maintaining high accuracy, making it well suited for real-time applications in ecological monitoring. These optimisations ensure that YOLOv10x performs well even in complex tasks involving multiple classes, such as wildlife monitoring for biodiversity assessments. Figure 4 shows the YOLOv10 architecture where numbered dots represent the detection proposals, with each number indicating a candidate bounding box for the object in the image.

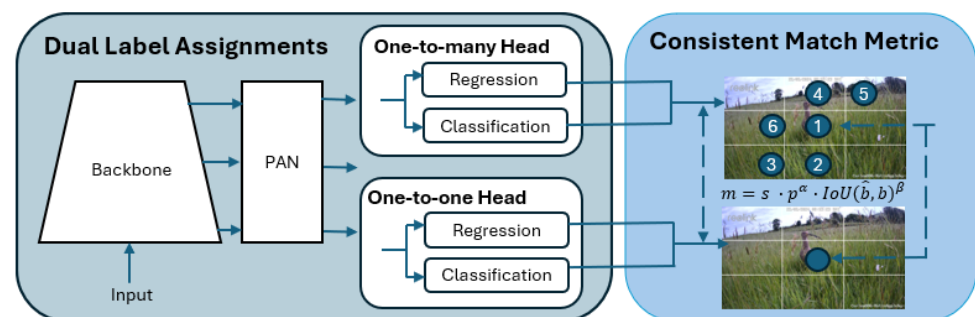


Figure 4. YOLOv10 architecture used in the study.

2.4. Transfer Learning

Transfer learning facilitates the adaptation of a pre-trained model to new tasks by fine-tuning its learned parameters for novel objects or species of interest [28]. This technique is critical when working with smaller datasets as training DL models from scratch on limited data can lead to poor performance due to reduced feature representation and low variance [29]. By leveraging a model pre-trained on a much larger dataset, transfer learning

helps mitigate these challenges, reducing the data required for effective training while maintaining robust accuracy. This makes it particularly valuable for applications with limited resources, such as ecological monitoring and biodiversity assessments.

In this study, YOLOv10x was employed as the foundation model for transfer learning. The model was initially pre-trained on the MS COCO dataset, a large object detection dataset allowing it to learn generalised features across a wide variety of objects. The pre-training of YOLOv10x was conducted using 8 NVIDIA 3090 GPUs over a span of approximately 120 h (500 epochs), as described in the original paper [30]. This setup enabled efficient optimisation of the model parameters, creating a strong foundation for transfer learning.

By fine-tuning YOLOv10x on a smaller dataset of 26 species, the model's learned weights were adapted to recognise new species with high accuracy, despite the limited size of the training data. This approach makes it feasible for a wide range of users to deploy high-performance models without requiring extensive computational resources or large datasets, making it an accessible solution for conservation and ecological studies.

2.5. Modelling

Model training was conducted on a custom-built Gigabyte server equipped with an AMD EPYC 7252 series processor and 128 GB of RAM. To support intensive computational tasks, the server was fitted with an additional GPU stack comprising 8 Nvidia Quadro A6000 graphics cards, which provided a combined total of 384 GB of GDDR6 memory. The training environment was set up using PyTorch 2.0.1 and CUDA 11.8, forming the core software components of the training pipeline. The following key hyperparameters were used during the training process:

- The image size is set to 640 pixels, ensuring the balance between detection accuracy and computational efficiency for high-resolution input. This resolution closely aligns with the mean resolution of the acquired dataset.
- The batch size coefficient is set to 256, allowing for a stable weight update process without exceeding the memory limitations of the available GPU hardware.
- The epoch parameter is configured to 50, ensuring adequate time for convergence while preventing overfitting in the model.
- The learning rate is set to 0.01, providing a balanced update speed, which prevents rapid shifts in response to errors.
- The momentum is set to 0.937, improving the training stability by maintaining model direction toward the minima during gradient descent.

To improve the generalisation of the model during inference and reduce overfitting, a number of different augmentation strategies were applied during training. It is important to note that these augmentations did not increase the size of the dataset as the images were randomly sampled and modified in real time during training. The applied augmentations included the following:

- Hue adjustment ($hsv_h = 0.015$): The hue of images was randomly adjusted by up to 1.5%, introducing slight colour variations.
- Saturation adjustment ($hsv_s = 0.7$): Saturation levels were altered by up to 70%, providing variety in colour intensity.
- Brightness adjustment ($hsv_v = 0.4$): The brightness (value) was adjusted by up to 40%, simulating different lighting conditions.
- Horizontal flip ($fliplr = 0.5$): Images were horizontally flipped with a 50% probability, increasing the model's invariance to directionality.
- Translation ($translate = 0.1$): Images were randomly shifted by up to 10%, helping the model handle variations in object positioning.

- Scaling (scale = 0.5): The size of objects in the images was adjusted by scaling up to 50%, improving detection at different object sizes.
- Random erasing (erasing = 0.4): Applied to 40% of the images, simulating partial occlusions by randomly removing parts of the image.

2.6. Model Inferencing

After training, the model was frozen and exported in ONNX format to facilitate inferencing. It was deployed and hosted via a publicly accessible website (www.conservationai.co.uk, accessed on 5 August 2024), developed by the authors for real-time species detection. The inference process was conducted on a custom-built server equipped with an Intel Xeon E5-1630v3 CPU, 256 GB of RAM, and an NVIDIA Quadro RTX 8000 GPU. The software stack for model inferencing utilised NVIDIA Triton Inference Server (version 22.08) [30], running within a Docker environment on Windows Subsystem for Linux 2 (WSL2). Given the high computational capabilities of the NVIDIA RTX 8000 GPU, no additional model optimisation techniques, such as quantisation or pruning, were necessary to enhance performance, ensuring efficient and accurate inferencing without compromising model complexity.

Next, 3/4G cellular cameras were deployed at 11 locations across the UK, as depicted in Figure 5. Each nesting site was monitored using a single camera; nine of the cameras were placed one meter away from the nest cup while the remaining two cameras were placed two meters away due to nest protection measures. The cameras were deployed in various habitats including two in heathland, eight in silage/hay, and one in rough grassland. The cameras were configured to capture images at a resolution of 1920×1072 pixels with a 96 DPI (dots per inch), closely matching the resolution of both the training images and the image size used during model training. The infrared (IR) sensor sensitivity was set to medium, and, when triggered, the acquired images were automatically uploaded to the platform for classification using the Simple Mail Transfer Protocol (SMTP). Each camera was powered by a lithium battery (7800 mAh Li9) and recharged via solar panels during their deployment. Figure 5 depicts two cameras deployed at different curlew monitoring locations, positioned in close proximity to curlew nests.



Figure 5. The 3/4G-enabled camera trap and solar panel deployed for real-time inference.

The end-to-end inferencing pipeline as shown in Figure 6 begins with data capture from the camera and concludes with the public-facing Conservation AI site. The system is designed to interface with a wide variety of cameras for real-time inference using standard protocols. When the IR sensor is triggered, the camera automatically transmits the image and associated meta data to a dedicated Simple Mail Transfer Protocol (SMTP) server running on the Conservation AI platform. The acquired data are automatically classified via the Triton Server REST API and saved to an internal database. In cases where field communication is unavailable, image files can be batch uploaded using the Conservation AI desktop application or programmatically using the REST API for offline inferencing, ensuring flexibility and scalability in data processing.




Figure 6. Conservation AI inferencing platform using Triton server.

The inferencing pipeline, illustrated in Figure 7, starts with the transmission of acquired image files and accompanying metadata (such as the Camera ID and image date/time) over 3G/4G networks using SMTP. The `imaplib` library is used to automatically extract images and associated payload data from the email body, which are then forwarded for classification on the Conservation AI platform. Upon receiving the image, the Triton inference server places the acquired data in a queue for processing. The classified image, along with probability scores, is logged in a MySQL database under the user's account if the prediction confidence exceeds 30%. Images with probability scores below 30% are categorised as blanks (no animal present) and are removed from the platform. Real-time alerts were configured to notify the conservation team when classification confidence exceeded 30% for specific classes, such as *Numenius arquata* (curlew) and *Numenius arquata* chick (curlew chick).

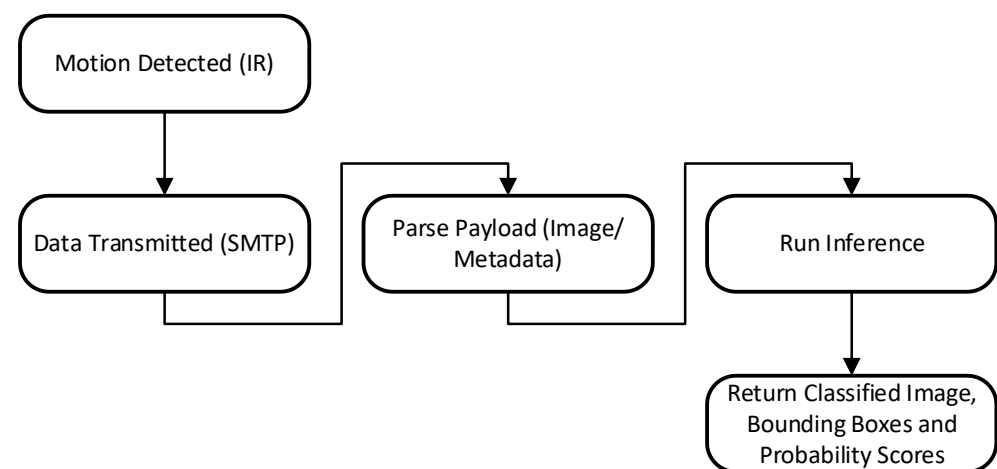


Figure 7. Conservation AI real-time inferencing pipeline.

2.7. Evaluation Metrics Training

The model is evaluated using the test split following training to assess its generalisation performance before deployment in real-time systems. This evaluation provides a comprehensive view of the model's behaviour on unseen data. To measure performance, several key metrics were employed: precision, recall (sensitivity), F1 score, and accuracy, along with visualisations such as the precision–confidence curve, recall–confidence curve, and F1–confidence curve.

The precision–confidence curve provides a visual representation of how the model's precision—its ability to make accurate positive detections—varies across different confidence thresholds. This curve helps in assessing the model's effectiveness in minimising false positives as the confidence level increases. A higher precision at higher confidence thresholds indicates that the model makes more accurate predictions when it is more certain about its detections.

The recall–confidence curve illustrates how the model's ability to correctly detect true positives changes across varying confidence thresholds. This curve is essential for understanding the model's sensitivity to false negatives. By examining the recall at different confidence levels, the curve provides insights into how well the model can capture all relevant objects, even as the confidence threshold increases.

The F1–confidence curve balances precision and recall at different confidence levels. This curve provides a more holistic view of the model's performance as the F1 score is a harmonic mean of both metrics. A consistently high F1 score across confidence thresholds indicates that the model is both accurate and reliable in detecting and classifying objects, with minimal trade-off between precision and recall.

2.8. Inference Evaluation Metrics

The performance of the trained model during inference is evaluated by analysing images transmitted from real-time cameras during the trial. Post-training analysis is crucial as model performance can vary significantly when deployed for real-world tasks due to variance not captured in the training data. Precision, sensitivity (recall), specificity, F1-score, and accuracy were employed to assess the model's performance on inference data. These metrics were derived from true positives (*TPs*), false positives (*FPs*), true negatives (*TNs*), and false negatives (*FNs*). Precision, for instance, is defined as follows:

$$\text{Precision} = \frac{TP}{TP + FP} \quad (1)$$

Precision measures the proportion of true-positive detections among all positive predictions made by the model. In the context of this study, it indicates how often the animal detected and classified by the model matches the ground truth, reflecting its accuracy in avoiding misclassifications. Recall, on the other hand, is defined as follows:

$$\text{Recall} = \frac{TP}{TP + FN} \quad (2)$$

Recall measures the proportion of true positives correctly identified by the model. In the context of object detection, it evaluates how effectively the model detects and correctly matches the ground-truth labels. Recall also helps identify the number of false negatives (instances the model missed). The F1-score, which balances precision and recall, is defined as follows:

$$F - 1 \text{ Score} = 2 * \frac{\text{Precision} * \text{Sensitivity (Recall)}}{\text{Precision} + \text{Sensitivity (Recall)}} \quad (3)$$

The F1-score represents the harmonic mean of precision and recall, providing a single metric that balances both. A high F1-score indicates that the model achieves a good trade-off between precision and recall. In the context of object detection, high recall signifies the model's ability to detect and localise most objects in an image, while precision ensures that the detections are accurate. Lastly, accuracy is defined as follows:

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \quad (4)$$

Accuracy provides an overall evaluation of the model's performance in detecting and correctly classifying objects within an image. However, its relevance diminishes in the context of unbalanced datasets as it can be misleading when certain classes are over- or under-represented. Therefore, accuracy should always be considered alongside other metrics, such as precision, recall, and the F1-score, to provide a more comprehensive assessment of model performance.

3. Results

The results are presented in two sections: first, the training of the YOLOv10x model using the hyperparameters and data detailed in the methodology; second, the inference results, showcasing the model's performance in a real-world setting using 3/4G real-time cameras to detect curlew and curlew chicks. It is important to note that the term "classes" refers to individual species or objects, and, as such, the terms "species" and "classes" are used interchangeably throughout this paper.

3.1. Training Results for UK Mammals Model

The tagged dataset was randomly divided into three subsets: 80% for the training set to train the model, 10% for the validation set to fine-tune and optimise the model, and 10% for the test set to evaluate the model's performance on unseen data. The division was undertaken on the tagged objects and not the number of images to ensure fair representation of each class. The model was trained over 50 epochs using a batch size of 256 to determine the best fit. During training, there was no overlap between the training and validation loss, indicating that no overfitting occurred.

The precision–recall (PR) curve presented in Figure 8 provides a detailed assessment of the model's performance across all 26 classes. The model achieved a mean average precision (mAP) of 0.976 at a 0.5 Intersection over Union (IoU) threshold, indicating a high level of overall detection accuracy. This high mAP value reflects the model's strong capability in both detecting and correctly classifying species across diverse categories. The curve demonstrates that the model sustains high precision even as recall increases, which is a clear indicator of the model's robustness and reliability in minimising false positives while capturing true positives. The shape of the PR curve remains stable across most classes, suggesting that the model is well calibrated for a wide range of detection tasks. However, some deviations are observed in the curves for individual classes, indicating that there are certain species for which the model's performance could be improved. These deviations could be attributed to insufficient training data for specific classes or challenges in distinguishing between visually similar species. Addressing these issues through further fine-tuning or the inclusion of additional training data may enhance the model's ability to generalise across all species.

The precision–confidence curve for the model, as illustrated in Figure 9, provides a detailed analysis of the model's reliability across all 26 classes. The steep ascent of the curve indicates that the model achieves high precision even at relatively low confidence

thresholds. This suggests that the model can make accurate predictions with moderate confidence, maintaining strong performance across a broad range of confidence levels.




Figure 8. Precision–recall (PR) curve for the model trained on UK species and objects.




Figure 9. Precision–confidence curve for the model trained on UK species and objects.

At the upper end of the confidence spectrum, the model achieves perfect precision at a confidence level of 1.0 (with a precision value of 1.00 at 0.996). This outcome demonstrates

that the model is consistently accurate when it is highly confident in its predictions. Minor variations observed in the curves for individual classes reflect some inherent classification challenges, likely due to interspecies similarities or variance in visual features. However, the overall correlation between precision and confidence confirms the robustness of the model in making reliable detections, even when confidence levels are not maximised. This consistent performance across confidence levels makes the YOLOv10x model highly suitable for real-world deployment as it can confidently handle predictions with both low and high certainty, maintaining accuracy and minimising false positives across varying conditions. These characteristics underscore the model's effectiveness in tasks requiring precision, such as wildlife species classification and biodiversity monitoring.

The recall–confidence curve, as presented in Figure 10, illustrates the relationship between recall and confidence across all 26 classes. At lower confidence thresholds, the model achieves near-perfect recall, with a recall value of 0.98 at a confidence threshold of 0.0. This indicates that the model is capable of capturing nearly all true positives when the confidence requirement is minimal, ensuring comprehensive detection coverage. As the confidence threshold increases, recall decreases gradually, with a more pronounced drop occurring near the highest confidence levels. This behaviour reflects the trade-off between recall and precision: while the model becomes more confident in its predictions at higher thresholds, it detects fewer objects, thus reducing recall. The sharp decline in recall at higher confidence thresholds suggests that fewer positive instances are identified when the model prioritises high certainty, potentially missing some true positives. The variation in individual class curves indicates that certain species may require lower confidence thresholds to maintain high recall. This underscores the need for careful tuning of the confidence threshold depending on the class or task. Overall, the model exhibits strong recall performance at lower confidence levels, which ensures that it captures a wide range of objects, making it suitable for applications where detecting all possible instances is critical.




Figure 10. Recall–confidence curve for the model trained on UK species and objects.

The F1–confidence curve in Figure 11 offers a detailed analysis of the model’s trade-off between precision and recall across varying confidence thresholds. The model achieves its peak F1-score of 0.96 at a confidence threshold of 0.387, which indicates that the model performs optimally at this threshold, striking an ideal balance between precision and recall. This high F1-score across a range of confidence levels demonstrates the model’s robustness in balancing these two key metrics for object detection tasks. As the confidence threshold increases towards 1.0, there is a noticeable decline in the F1-score. This sharp drop suggests that, while the model becomes more confident in its predictions, it begins to miss more true positives, leading to a decrease in recall and a consequent decline in the F1-score. This behaviour reflects the inherent trade-off between precision and recall: as precision increases, fewer true positives are detected, leading to a drop in overall performance as measured by the F1-score. Despite this, the model maintains a consistently high F1-score across most confidence levels, underscoring its suitability for tasks that require a delicate balance between precision and recall. The model’s performance at these varying confidence thresholds indicates its effectiveness in handling complex detection scenarios, making it particularly well suited for applications in species detection and classification within ecological monitoring frameworks.




Figure 11. F1–confidence curve for the model trained on UK species and objects.

The confusion matrix in Figure 12 provides a comprehensive assessment of the model’s classification performance across the 26 species, vehicles, human subjects, and background categories used in this study. Each row corresponds to the predicted class, while each column represents the true class. The strong diagonal line seen in the matrix reflects that the majority of predictions are correct, with high values along the diagonal indicating accurate classification for most species. Notably, the most frequent species in the dataset, such as European hedgehog, eastern grey squirrel, and roe deer, exhibit the highest intensities along the diagonal, suggesting that the model has effectively learned to identify these species with high precision. The matrix also demonstrates the model’s ability to distinguish between visually distinct species. However, several off-diagonal cells with lighter shading indicate instances of misclassification. These misclassifications, although

relatively infrequent, occur between species that may share similar visual traits or appear in complex environmental conditions. For example, there are occasional confusions between roe deer and red deer, likely due to similarities in their appearance or habitat. Similarly, the misclassification between house sparrow and other bird species could be attributed to the challenge of distinguishing between small birds in various environmental conditions. Overall, the confusion matrix illustrates the model’s strong classification performance, with high accuracy across the majority of classes and minimal instances of misclassification. The results highlight the model’s robustness in distinguishing between a diverse range of species, supporting its applicability in ecological monitoring and species classification tasks.




Figure 12. Confusion matrix for the model trained on UK species and objects.

3.2. Model Deployment

The trained model was deployed to evaluate its performance during the trial period. Using the inferencing pipeline and camera setup described in the Methodology section, the system was deployed to monitor nesting curlews at 11 locations in Wales, UK. This deployment spanned from 20 May 2024 to 30 June 2024, during which a total of 1072 images were analysed by the platform. These images contained three distinct detections: domestic sheep (detected as *Ovis aries*), common curlew (detected as *Numenius arquata*) and common curlew (chick) (detected as *Numenius arquata* chick). The detected objects in these images were evaluated to assess the model’s detection accuracy and classification performance for curlews and their chicks, as well as for other species detected during the study. This analysis was critical for understanding the model’s effectiveness in real-world scenarios and its suitability for long-term monitoring of species.

3.2.1. Performance Evaluation Results for Inference

The model achieved a high level of accuracy across most classes, with individual class accuracies ranging from 93.41% to 100% and an overall accuracy of 91.23% (Table 1). However, the performance metrics exhibited considerable variation among classes. Precision

ranged from 0% to 100%, with common pheasant (detected as *Phasianus colchicus*) having the lowest value due to misclassifications and the absence of true instances in the dataset, while common curlew (adult) and domestic sheep achieved perfect scores of 100%. The average precision of 60.34% suggests that while the model performs well for certain classes, there are significant issues with others.

Table 1. Inference performance metrics for camera deployment.

	Accuracy	Precision	Sensitivity	Specificity	F1-Score
<i>Numenius arquata</i>	93.41%	100%	90.56%	100%	95.05%
<i>Numenius arquata</i> chick	97.51%	100%	92.35%	100%	96.03%
<i>Ovis aries</i>	100%	100%	100%	100%	100%

Notably, some classes were discontinued from the analysis due to the absence of true instances observed during the study. Specifically, common pheasant was not present in the dataset; however, some common curlews (adult) were misclassified as common pheasant. This misclassification led to false positives for common pheasant, resulting in a precision and F1-score of 0% for this class. The sensitivity of the model was high across classes with actual instances, ranging from 90.56% for common curlew (adult) to 100% for domestic sheep. The average sensitivity of 95.48% indicates that the model can correctly identify a large proportion of actual positive cases. Specificity was also notably high, with values between 96.64% and 100%. The species with the lowest specificity was common pheasant, reflecting the misclassification issues, while common curlew (adult) and domestic sheep each attained a specificity of 100%. The model's average specificity of 98.17% demonstrates a strong ability to correctly identify negative cases. The F1-scores varied widely across classes, from 0% for common pheasant to 100% for domestic sheep. The species with the highest F1-score was domestic sheep (100%), while the lowest was common pheasant (0%). The average F1-score across all classes was 58.88%, indicating that while the model performs exceptionally for some classes, it underperforms for others. The particularly low F1-score for common pheasant underscores the impact of misclassification and the presence of classes without true instances in the dataset. This highlights the need for careful dataset curation and potential refinement of the model to address misclassification issues.

3.2.2. Confusion Matrix for Inference Data

Sources of confusion within the inference data (Table 1) closely correspond to the performance metrics outlined in Table 2. The model demonstrates strong predictive accuracy, correctly classifying most classes with minimal error; the common curlew (adult) and domestic sheep classes performed the best. However, misclassifications were observed, notably with some instances of common curlew (adult) being incorrectly classified as common pheasant despite the latter not being present in the actual dataset. This misclassification suggests that the model may rely on features common between these species, leading to false positives. These results highlight weaknesses in the model's ability to accurately detect and classify certain species, which will need to be addressed in future training iterations. Improvements may include refining the training dataset, enhancing feature selection, and adjusting augmentation techniques to reduce misclassification rates.

Table 2. Confusion matrix for inference. The diagonal number indicates the TP for each of the classes.

Actual \ Predicted	<i>Numenius arquata</i>	<i>Numenius arquata</i> chick	<i>Phasianus colchicus</i>	<i>Ovis aries</i>	Total Actual
<i>Numenius arquata</i>	662	0	36	33	731
<i>Numenius arquata</i> chick	0	302	0	25	327
<i>Ovis aries</i>	0	0	0	13	13
<i>Phasianus colchicus</i>	0	0	0	0	0
Total Predicted	662	302	49	58	1072

4. Discussion

This study demonstrates the effectiveness of an AI-driven approach for real-time monitoring of ground-nesting birds, with a focus on curlew (*Numenius arquata*) detection. Leveraging the YOLOv10x architecture, the model was trained on a diverse dataset, achieving high classification performance for curlews and their chicks. Integration with the Conservation AI platform allowed the model to be deployed in conservation study pipelines to process data in real time and to deliver actionable insights.

During the evaluation phase, the model processed 1072 images from 11 camera sites in Wales. The results showed an overall accuracy of 91.23%, with specificity reaching 98.17% and sensitivity at 95.48%. For curlew detections, the model achieved a precision score of 100%, reflecting its ability to accurately identify true positives without misclassifications. Similarly, the F1-scores for adult curlews and chicks was 95.05% and 96.03%, respectively, underscoring the model's robust performance in both precision and recall.

The model's performance remained consistent across diverse and challenging environmental conditions. Figures 13 and 14 showcase successful detections of adult curlews and chicks, even in complex backgrounds. This is particularly significant for smaller species like curlew chicks, which are prone to blending into their natural surroundings.



Figure 13. Example detections of common curlew (adult) from one of the real-time cameras classified by the Conservation AI platform.

As illustrated in Figure 15, the model effectively distinguished curlews from dense vegetation, shadows, and other environmental noise, demonstrating its ability to generalise across variable conditions. This ability is critically important given the nature of the nesting locations of curlews.



Figure 14. Example detection of a common curlew chick from one of the real-time cameras classified by the Conservation AI platform.



Figure 15. Example detection of a curlew and curlew chick obtained from a real-time camera.

The integration of the YOLOv10x model into a real-time monitoring system addresses critical limitations in traditional conservation methods. By filtering out irrelevant images with an accuracy of 98.28%, the system significantly reduces the data processing burden on conservationists, enabling them to focus on relevant detections and respond more promptly to potential threats. Additionally, the automated detection pipeline mitigates delays inherent in manual data processing, providing an efficient and scalable solution for biodiversity monitoring. This enhanced responsiveness allows conservationists and researchers to intervene much quicker when alerts are raised, maximising the protection of curlew eggs and chicks against predation and other threats.

Despite these achievements, several issues were identified. Misclassifications, particularly between common pheasants and curlews, were notable during the study. Furthermore, correct camera placement was a critical factor for detecting smaller subjects like curlew chicks. Camera traps need to be deployed closer to nesting sites to ensure optimal trigger settings and to capture high-quality images in dynamic or cluttered environments, like the grassy locations in which curlews' nest. Not all camera trap installations in the study

adhered to these guidelines; consequently, some misdetections were observed. In addition to the misclassification of curlews as pheasants, there were instances where curlews were incorrectly identified as sheep. We speculate that this error may arise from overlapping visual features, such as similar colouration under certain lighting conditions or partial captures of the birds that obscure distinctive characteristics. Furthermore, environmental factors like rain or condensation on the camera lens can impair image clarity, making it challenging for the model to accurately discern specific features of curlews.

It is important to acknowledge that the detection of species not represented in the model's training set can lead to misclassifications. Therefore, expanding the training dataset to include a broader spectrum of species is necessary to improve the model's accuracy and reliability. The cameras were deployed in slightly varying habitats although the overall colour and structure of the vegetation did not vary significantly between each site. However, the density of vegetation especially in the hay and silage fields varied depending on the growth. Model performance would likely degrade in environments that differ significantly from the habits used in this study. It would, therefore, be important to add additional training data to increase environmental variance.

5. Conclusions

This study introduced an AI-driven classification system that utilises real-time 3/4G-enabled cameras to detect and monitor curlew populations. By automating image collection and classification, the system overcomes the challenges of large-scale monitoring in remote areas, reducing reliance on manual processing. This approach accelerates data processing, providing timely information that allows conservationists to respond swiftly to threats such as habitat loss and predation.

This paper presented a custom-built pipeline integrating real-time 3/4G-enabled camera traps with a YOLOv10x model specifically trained for curlew detection. This system leverages cutting-edge sensor and AI technology, enabling immediate insights for conservation interventions, which would not be possible with traditional wildlife monitoring approaches. Using the YOLOv10x architecture and transfer learning, the model was deployed across 11 sites in Wales to evaluate its effectiveness in monitoring curlews and their chicks. Real-time data transmission facilitated immediate classification and documentation of curlew activity, showcasing the system's potential for proactive conservation efforts. For instance, this system could support breeding curlew recovery initiatives by enabling the early detection of threats or environmental stressors, informing targeted interventions to protect vulnerable nests and improve chick survival rates. This contribution to curlew population recovery is critical for reversing the alarming declines experienced by this species.

The results of the trial were promising, demonstrating high accuracy in curlew detection. The model effectively filtered out blank images triggered by moving vegetation, significantly reducing the workload for conservationists. Hosting the model on the Conservation AI platform further enhanced its scalability, allowing for broader deployment to monitor other species or to expand curlew conservation efforts.

This automated monitoring system also creates opportunities for broader community engagement. By encouraging individuals to deploy cameras in local environments—such as gardens, communal spaces, or workplaces—this approach could greatly expand the reach of conservation initiatives for other types of nesting birds—not just ground-nesting species. Organisations like the Game & Wildlife Conservation Trust (GWCT) can now leverage this scalable solution to enhance biodiversity monitoring, enabling a more comprehensive understanding of species distribution and population dynamics across the UK. Engaging citizen scientists in this way not only increases data collection capacity but also fosters

public awareness and involvement in conservation efforts. Furthermore, citizen scientists can play a pivotal role in model development by contributing additional training data and assisting in data annotation. This collaborative effort enhances the accuracy and robustness of the AI-driven classification system, facilitating continuous improvement and adaptation to diverse environmental conditions and species-specific characteristics. However, the integration of citizen science also presents significant challenges. The absence of expert annotation can lead to inconsistencies and potential inaccuracies in the data, as non-expert participants may misidentify species or incorrectly annotate images. Additionally, data quality issues, such as variability in camera placement, lighting conditions, and image resolution, can affect the reliability of the collected data. Addressing these challenges requires the implementation of comprehensive training programs for citizen scientists, the development of standardised protocols for data collection and annotation, and the incorporation of robust data validation mechanisms to ensure the integrity and accuracy of the dataset.

Overcoming the issues highlighted in Section 4, future work will focus on refining model performance, particularly through guidelines for camera placement and classification accuracy of smaller species like curlew chicks. Enhancements in training data quality, data augmentation techniques, and hyperparameter tuning are also needed to train out misclassification such as those evident between common pheasants and curlew.

Overall, the study demonstrated excellent utility in monitoring ground-nesting birds. The approach provides a scalable and effective tool to support curlew conservation and recovery efforts, aligning with both immediate and long-term biodiversity goals. By making the model publicly available on the Conservation AI platform, this research empowers researchers, conservationists, and citizen scientists to contribute meaningfully to curlew population recovery and broader conservation initiatives.

This work had provided us with a solid platform for the longitudinal curlew nesting season survey in 2025, allowing us to evaluate the efficiency of our results over time and assess the potential impact of this approach on curlew populations.

Author Contributions: Conceptualisation, C.C. and P.F.; methodology, C.C., P.F., S.N.L., S.W.; software, C.C. and P.F.; validation, C.C., K.A., L.O., J.W. and J.Q.; formal analysis, C.C. and P.F.; investigation, C.C., K.A., L.O., J.W., J.Q. and N.D.W.; resources, C.C. and P.F.; data curation, C.C., K.A. and J.Q.; writing—original draft preparation, C.C., P.F., S.N.L., S.W., N.D.W., K.A., L.O., J.W. and J.Q.; writing—review and editing, C.C. and P.F.; visualisation, C.C. and P.F.; supervision, C.C.; project administration, C.C.; All authors have read and agreed to the published version of the manuscript.

Funding: The research received no external funding.

Data Availability Statement: Data is unavailable due privacy and prior data sharing agreements.

Acknowledgments: The authors would like to thank ReoLink for donating the cameras and solar panels used in this study and Vodafone UK for sponsoring our communication. Finally, the authors would like to thank Rachel Chalmers for tagging all the data used in this study and for contributing so much to the Conservation AI platform.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Burns, F.; Mordue, S.; al Fulaij, N.; Boersch-Supan, P.H.; Boswell, J.; Boyd, R.J.; Bradfer-Lawrence, T.; de Ornellas, P.; de Palma, A.; de Zylva, P.; et al. State of nature. *State Nat. Partnersh.* **2023**, *2023*, 213.
2. Tobajas, J.; Descalzo, E.; Mateo, R.; Ferreras, P. Reducing nest predation of ground-nesting birds through conditioned food aversion. *Biol. Conserv.* **2020**, *242*, 108405. [[CrossRef](#)]
3. Roodbergen, M.; van der Werf, B.; Hoetker, H. Revealing the contributions of reproduction and survival to the Europe-wide decline in meadow birds: Review and meta-analysis. *J. Ornithol.* **2012**, *153*, 53–74. [[CrossRef](#)]

4. Bravo, C.; Pays, O.; Sarasa, M.; Bretagnolle, V. Revisiting an old question: Which predators eat eggs of ground-nesting birds in farmland landscapes? *Sci. Total Environ.* **2020**, *744*, 140895. [[CrossRef](#)] [[PubMed](#)]
5. Whelan, C.J.; Wenny, D.G.; Marquis, R.J. Ecosystem Services Provided by Birds. *Ann. N. Y. Acad. Sci.* **2008**, *1134*, 25–60. [[CrossRef](#)]
6. Díaz-Sieffer, P.; Olmos-Moya, N.; Fontúrbel, F.E.; Lavandero, B.; Pozo, R.A.; Celis-Diez, J.L. Bird-mediated effects of pest control services on crop productivity: A global synthesis. *J. Pest. Sci.* **2022**, *95*, 567–576. [[CrossRef](#)]
7. Bocher, P.; Donnez, M.; Chenu, A.; Sviridova, T.; Fort, J.; Garthe, S.; Jiguet, F.; Piha, M.; Elts, J.; Marja, R.; et al. Home ranges and hatching success of threatened Eurasian curlew in north-eastern Europe relates to habitat type: Natural vs. agricultural landscapes. *Glob. Ecol. Conserv.* **2024**, *50*, e02851. [[CrossRef](#)]
8. Viana, D.S.; Santoro, S.; Soriguer, R.C.; Figuerola, J. A synthesis of Eurasian Curlew (*Numenius arquata arquata*) demography and population viability to inform its management. *Ibis* **2023**, *165*, 767–780. [[CrossRef](#)]
9. Ewing, H.; Franks, S.; Smart, J.; Burton, N.; Gill, J.A. Nest survival of threatened Eurasian Curlew (*Numenius arquata*) breeding at low densities across a human-modified landscape. *Ibis* **2023**, *165*, 753–766. [[CrossRef](#)]
10. Gautschi, D.; Čulina, A.; Heinsohn, R.; Stojanovic, D.; Crates, R. Protecting wild bird nests against predators: A systematic review and meta-analysis of non-lethal methods. *J. Appl. Ecol.* **2024**, *61*, 1187–1198. [[CrossRef](#)]
11. O'Brien, T.G. Abundance, Density and Relative Abundance: A Conceptual Framework. In *Camera Traps in Animal Ecology*; Springer: Tokyo, Japan, 2011; pp. 71–96. [[CrossRef](#)]
12. Wearn, O.R.; Glover-Kapfer, P. Snap happy: Camera traps are an effective sampling tool when compared with alternative methods. *R. Soc. Open Sci.* **2019**, *6*, 181748. [[CrossRef](#)] [[PubMed](#)]
13. Young, S.; Rode-Margono, J.; Amin, R. Software to facilitate and streamline camera trap data management: A review. *Ecol. Evol.* **2018**, *8*, 9947–9957. [[CrossRef](#)] [[PubMed](#)]
14. Chalmers, C.; Fergus, P.; Wich, S.; Longmore, S.N.; Walsh, N.D.; Stephens, P.A.; Sutherland, C.; Matthews, N.; Mudde, J.; Nuseibeh, A. Removing Human Bottlenecks in Bird Classification Using Camera Trap Images and Deep Learning. *Remote Sens.* **2023**, *15*, 2638. [[CrossRef](#)]
15. Vélez, J.; McShea, W.; Shamon, H.; Castiblanco-Camacho, P.J.; Tabak, M.A.; Chalmers, C.; Fergus, P.; Fieberg, J. An evaluation of platforms for processing camera-trap data using artificial intelligence. *Methods Ecol. Evol.* **2023**, *14*, 459–477. [[CrossRef](#)]
16. Norouzzadeh, M.S.; Nguyen, A.; Kosmala, M.; Swanson, A.; Palmer, M.S.; Packer, C.; Clune, J. Automatically identifying, counting, and describing wild animals in camera-trap images with deep learning. *Proc. Natl. Acad. Sci. USA* **2018**, *115*, E5716–E5725. [[CrossRef](#)]
17. Steenweg, R.; Hebblewhite, M.; Kays, R.; Ahumada, J.; Fisher, J.T.; Burton, C.; E Townsend, S.; Carbone, C.; Rowcliffe, J.M.; Whittington, J.; et al. Scaling-up camera traps: Monitoring the planet's biodiversity with networks of remote sensors. *Front. Ecol. Environ.* **2017**, *15*, 26–34. [[CrossRef](#)]
18. Beery, S.; Morris, D.; Yang, S. Efficient pipeline for camera trap image review. *arXiv* **2019**, arXiv:1907.06772.
19. Hernandez, A.; Miao, Z.; Vargas, L.; Dodhia, R.; Lavista, J. Pytorch-Wildlife: A Collaborative Deep Learning Framework for Conservation. *arXiv* **2024**, arXiv:2405.12930.
20. Ahumada, J.; Fegraus, E.; Birch, T.; Flores, N.; Kays, R.; O'Brien, T.G.; Palmer, J.; Schuttler, S.; Zhao, J.Y.; Jetz, W.; et al. Wildlife Insights: A Platform to Maximize the Potential of Camera Trap and Other Passive Sensor Wildlife Data for the Planet. *Environ. Conserv.* **2020**, *47*, 1–6. [[CrossRef](#)]
21. Van Horn, G.; Mac Aodha, O.; Song, Y.; Cui, Y.; Sun, C.; Shepard, A.; Adam, H.; Perona, P.; Belongie, S. The inaturalist species classification and detection dataset. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Salt Lake City, UT, USA, 18–22 June 2018; pp. 8769–8778.
22. Swanson, A.; Kosmala, M.; Lintott, C.; Simpson, R.; Smith, A.; Packer, C.; Serengeti, S. high-frequency annotated camera trap images of 40 mammalian species in an African savanna. *Sci. Data* **2015**, *2*, 150026. [[CrossRef](#)] [[PubMed](#)]
23. Hu, L.; Robinson, C.; Dilkina, B. Model generalization in deep learning applications for land cover mapping. *arXiv* **2020**, arXiv:2008.10351.
24. Pahuja, V.; Luo, W.; Gu, Y.; Tu, C.H.; Chen, H.Y.; Berger-Wolf, T.; Stewart, C. Bringing back the context: Camera trap species identification as link prediction on multimodal knowledge graphs. *arXiv* **2023**, arXiv:2401.00608.
25. Wong, W.-M.; Kachel, S. *Camera Trapping—Advancing the Technology*; Leopards, S., Ed.; Elsevier: Amsterdam, The Netherlands, 2024; pp. 415–428. [[CrossRef](#)]
26. Fergus, P.; Chalmers, C.; Longmore, S.; Wich, S. Harnessing Artificial Intelligence for Wildlife Conservation. *arXiv* **2024**, arXiv:2409.10523. [[CrossRef](#)]
27. Wang, A.; Chen, H.; Liu, L.; Chen, K.; Lin, Z.; Han, J. Yolov10: Real-time end-to-end object detection. *arXiv* **2024**, arXiv:2405.14458.
28. Cui, Y.; Song, Y.; Sun, C.; Howard, A.; Belongie, S. Large scale fine-grained categorization and domain-specific transfer learning. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Salt Lake City, UT, USA, 18–22 June 2018; pp. 4109–4118.

29. Romero, M.; Interian, Y.; Solberg, T.; Valdes, G. Targeted transfer learning to improve performance in small medical physics datasets. *Med. Phys.* **2020**, *47*, 6246–6256. [[CrossRef](#)] [[PubMed](#)]
30. Savard, C.; Manganelli, N.; Holzman, B.; Gray, L.; Perloff, A.; Pedro, K.; Stenson, K.; Ulmer, K. Optimizing High-Throughput Inference on Graph Neural Networks at Shared Computing Facilities with the NVIDIA Triton Inference Server. *Comput. Softw. Big Sci.* **2024**, *8*, 14. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.