

Long-tailed Species Recognition in the NACTI Wildlife Dataset

Zehua Liu¹ ^a and Tilo Burghardt¹ ^b

¹*School of Computer Science, University of Bristol, MVB Woodland Rd, BS8 1UB, Bristol, UK
[bw19062@bristol.ac.uk; tilo@cs.bris.ac.uk]*

Keywords: Long-tail Classification, Animal Biometrics, Computer Vision, AI for Ecology

Abstract: As most “in the wild” data collections of the natural world, the North America Camera Trap Images (NACTI) dataset shows severe long-tailed class imbalance, noting that the largest ‘Head’ class alone covers >50% of the 3.7M images in the corpus. Building on the PyTorch Wildlife model, we present a systematic study of Long-Tail Recognition methodologies for species recognition on the NACTI dataset covering experiments on various LTR loss functions plus LTR-sensitive regularisation. Our best configuration achieves 99.40% Top-1 accuracy on our NACTI test data split, substantially improving over a 95.51% baseline using standard cross-entropy with Adam. This also improves on previously reported top performance in MLWIC2 at 96.8% albeit using partly unpublished (potentially different) partitioning, optimiser, and evaluation protocols. To evaluate domain shifts (e.g. night-time captures, occlusion, motion-blur) towards other datasets we construct a Reduced-Bias Test set from the ENA-Detection dataset where our experimentally optimised long-tail enhanced model achieves leading 52.55% accuracy (up from 51.20% with WCE loss), demonstrating stronger generalisation capabilities under distribution shift. We document the consistent improvements of LTR-enhancing scheduler choices in this NACTI wildlife domain, particularly when in tandem with state-of-the-art LTR losses. We finally discuss qualitative and quantitative shortcomings that LTR methods cannot sufficiently address, including catastrophic breakdown for ‘Tail’ classes under severe domain shift. For maximum reproducibility we publish all dataset splits, key code, and full network weights.




Figure 1: Representative Samples from the NACTI Dataset. Shown is one representative image per species class from the overall 48 classes in total, excluding empty and vehicle frames. Note challenging occlusions, background clutter, motion blur, infrared shots, and small animal areas. Red bounding boxes indicate detections using MegaDetectorV6, applied during the pre-processing stage. The NACTI dataset is long-tailed with over half of the samples belong to a single species (domestic cow). Following (Perrett et al., 2023), 5 classes form the head, 37 form the ‘long tail’, and 6 few-shot classes sit at the end of the species distribution.

^a  <https://orcid.org/0009-0004-5353-7028>

^b  <https://orcid.org/0000-0002-8506-012X>

1 INTRODUCTION

The Challenging NACTI Wildlife Dataset. The North American Camera Trap Images (NACTI) dataset (LILA Science, 2024b) – visually exemplified in Figure 1 – contains 3.7M images from five U.S. sites covering 48 animal (sub)species. It presents a formidable challenge for automated wildlife recognition (Kühl and Burghardt, 2013; Reynolds et al., 2025) due to extreme class imbalances. In NACTI, over half of the samples belong to a single species (domestic cow). Following (Perrett et al., 2023), 5 classes form the head, 37 form the ‘long tail’, and 6 few-shot classes sit at the end of the distribution where rare or endangered species are represented by often only a handful of examples. Furthermore, the dataset’s untrimmed images captured under uncontrolled conditions (e.g., low illumination, motion blur) and its lack of bounding box or auxiliary meta-data make fine-grained species ID challenging.

Limitations of Standard Vision Approaches. Discussed issues are characteristic of most collected data from natural settings that conservationists and ecologists work with on a daily basis. While Convolutional Neural Networks (CNNs) as well as Transformer-based architectures can often achieve near human-level accuracy on images taken under ideal conditions (Krizhevsky et al., 2012) and (Dosovitskiy et al., 2020), their performance degrades sharply on long-tailed natural data. Tools like MegaDetector (Beery et al., 2019) and SpeciesNet (Gadot et al., 2024) trained on many millions of images offer a strong foundation for general animal detection, but do not fully resolve the core issue of species-level classification under truly severe class imbalance. Consequently, standard models often fail decisively on minority classes, compromising the recognition of ecological animal encounters they are meant to provide for rare or endangered species that form the distribution tail (Chen et al., 2014).

Our Approach and Contributions. Towards tackling this critical issue in computer vision for wildlife monitoring (Tuia et al., 2022), our work presents a systematic, loss-centric study focused on improving minority class detection performance in the large-scale, real-world NACTI setting. We develop and present a scalable, multi-GPU pipeline to fine-tune state-of-the-art CNN architectures on the full NACTI corpus. Using this pipeline, we conduct a head-to-head comparison of Long-Tail Recognition (LTR) scheduling (Alshammari et al., 2022) and losses: focal loss (Lin et al., 2017), weighted cross-entropy (Cui et al., 2019), and label-distribution-aware margin (LDAM) loss (Cao et al., 2019). Re-

sulting from this analysis, Figure 2 summarises best model accuracy improvements over the baseline at class level. We also assess the generalisation potential of this model on a second camera-trap collection to quantify its robustness under distributional shift. Our paper contributions are threefold:

1. **First Systematic NACTI LTR Evaluation.** The first head-to-head evaluation of LTR-sensitive scheduling in tandem with focal, weighted cross-entropy, and LDAM LTR losses for large-scale wildlife classification, analysing trade-offs between minority-class recall and overall accuracy;
2. **Additional Out-of-Distribution (OOD) Robustness Analysis.** A cross-dataset assessment demonstrating that the best-performing loss maintains superior minority-class F_1 when species frequencies and imaging conditions change; and
3. **Reproducible Best Performance Classifiers.** Given limitations regarding the disclosure of details in much of prior art, we publish reproducible, fully open access multi-GPU networks that set the state-of-the-art on NACTI end-to-end and support easily comparable community research in future. We publish splits, code, and describe the approach in detail.

2 KEY RELATED WORK

2.1 Species Recognition in NACTI

General ResNet-based Approaches. Prior art on species recognition in NACTI goes back several years, but most published works lack reproducibility and published descriptions/code that would allow for a fair and detailed comparison of approaches to illuminate the effect of chosen architectures regarding the long-tail distribution issue. Tabak et al. (Tabak et al., 2022), for instance, trained a ResNet-50 architecture on the NACTI dataset (70/30 split) using SGD with momentum – however, the paper does not explicitly specify the loss function used; though one would assume a softmax-cross-entropy setting was applied. They reported $\approx 86\%$ accuracy on the validation set. Performance dropped significantly for under-represented species (recall $< 70\%$), highlighting persistent long-tailed bias in camera-trap datasets.

Cross-Study MLWIC2 Project. Tabak et al. (Tabak et al., 2020) also trained MLWIC2, that is a ResNet-18 on $> 3\text{M}$ images from 18 North American studies (including NACTI), reporting 96.8% Top-1 accuracy for the species model. To mitigate imbalance, the authors capped each class at 100,000 images. Despite these measures, rare species still showed weaker recall, indicating unresolved long-tail




Figure 2: **LTR Performance Improvements across NACTI Species vs. Baseline.** Per-class Top-1 accuracies on the NACTI dataset comparing the baseline model (cross-entropy with Adam) against our best-performing configuration, which combines a Label-distribution-aware margin (LDAM) loss, an alternative AdamW optimiser, and a learning-rate scheduler. The results are reported on a 10% random test split. The blue bars indicate classes with a positive improvement relative to the baseline.

effects. For out-of-sample evaluations, reported Top-5 accuracy ranged from 65.2% to 93.8% depending on the dataset. We note, that the paper does neither specify optimiser nor loss function used.

Active Learning. Norouzzadeh et al. (Norouzzadeh et al., 2021) combined a ResNet-50 with active learning techniques to iteratively select informative, unlabelled images for annotation. After 30,000 queries, the classifier reached 93.2% accuracy on NACTI. In their pipeline, the embedding network is trained either using softmax-cross-entropy or with triplet loss settings. However, the exact evaluation split is not specified and the optimiser used for these experiments is not explicitly specified in the paper either (SGD is discussed only as general background), so we treat it as unspecified.

Domain-aware DANAS Approach. Jia et al. (Jia et al., 2022) introduced DANAS, a domain-aware neural architecture search tailored for camera-trap species classification. On the NACTI-a subset, DANAS showed that lightweight CNNs can match conventional performance baselines (e.g., ResNet-18) within Top-1/Top-5 accuracy charts while using far fewer parameters and lower compute, enabling edge deployment at average accuracy $\approx 92.9\%$. During the search phase, candidate networks were trained with AMSGrad (learning rate = 0.005, batch size = 32) and guided by a customised, theoretically derived losses (Witch-of-Agnesi-based) that encourage high accuracy with few parameters. The optimiser used for the final training of the discovered DANAS networks is not explicitly stated, while baselines were stated and trained with Nesterov-SGD under a cosine loss regime (Jia et al., 2022).

2.2 Regularisers and Losses for LTR

Scheduling for LTR. Addressing class imbalance, a central challenge in real-world data (Zhang et al., 2021), has led to various strategies. Recent study (Alshammari et al., 2022) have highlighted the importance of weight decay in mitigating classifier bias towards dominant classes, which is a key challenge in long-tailed recognition tasks. Alshammari et al. (Alshammari et al., 2022) showed that models trained without weight decay tend to assign disproportionately large classifier weights to frequent classes, resulting in poor generalisation for rare classes. By carefully tuning weight decay parameters, they significantly improved recognition performance on long-tailed datasets. Unlike conventional L1 or L2 regularisation, their method effectively balanced per-class classifier weights, thereby preventing common species from dominating learning. Furthermore, their study demonstrated that simple weight decay tuning outperformed many sophisticated LTR methods, suggesting that parameter regularisation should be of primary consideration when addressing class imbalance.

Decoupled Learning Strategies. A seminal work by Kang et al. (Kang et al., 2020) proposed decoupling feature representation learning from classifier learning. They observed that standard instance-balanced sampling yields high-quality representations, while the classifier benefits from a class-balanced perspective, e.g. re-weighting or re-sampling. Their work also introduced LDAM-DRW, which combines the LDAM loss with a *Deferred Re-weighting (DRW) schedule*, delaying the application of class-balancing until later in training. This concept strongly aligns with our findings, where combining




Figure 3: **NACTI Species Distribution after Sample Balancing.** Logarithmic visualisation of the cardinality of all 48 (sub)species classes of NACTI after capping classes at 100k. Colouration of the class distribution after (Perrett et al., 2023) breaking the distribution into ‘Head’ (blue), ‘Tail’ (pink), and ‘Few-shot’ species.

the LDAM loss with a learning rate (LR) scheduler (ReduceLROnPlateau) yielded best results (see Section 4.2 for details).

Problem-specific LTR Loss Concepts. To systematically address the severe class imbalance in the NACTI dataset, we implemented, compared, and evaluated three state-of-the-art loss functions that modify the standard cross-entropy in complementary ways. Our goal was to isolate the most effective strategy for improving performance on minority classes within a large-scale, real-world setting.

Focal Loss. Focal loss (Lin et al., 2017) reduces the weight of easily classified samples, encouraging the model to focus more on hard-to-classify examples. This improves performance in classes with fewer samples.

Weighted Cross-Entropy (WCE) Loss. The WCE concept (Cui et al., 2019) addresses class imbalance by assigning greater weights to minority classes, thereby encouraging the model to focus more on under-represented categories during training and improving their recognition rates. However, this can lead to overfitting, causing the model to prioritise minority classes at the expense of overall performance.

Label-Distribution-Aware Margin (LDAM) Loss. LDAM (Cao et al., 2019) addresses class imbalance by encouraging larger decision margins for minority classes. It modifies the standard cross-entropy loss by subtracting a class-dependent margin from the logit of the true class. This formulation gently penalises confident predictions for frequent classes while imposing stronger regularisation on minority classes. As a result, LDAM improves generalisation across the label distribution, especially for under-represented classes. It can also be combined with re-weighting or re-sampling strategies for further performance gains.

3 DATA PREPARATION

3.1 Dataset Statistics

The NACTI Dataset. Published in the LILA Dataset collection (LILA Science, 2024b), NACTI contains 3.7M camera-trap images from spanning 48 (sub)species. Figure 1 displays representative images from each species class, illustrating the wide variability in appearance, lighting, and environmental conditions. These samples exemplify common visual challenges encountered in camera trap imagery, including occlusions, background clutter, motion blur, and significant class imbalance. Such factors complicate the training process significantly particularly for under-represented classes where few images are available to capture species-distinctive features.

Species Distribution Details. The class distribution is extremely long-tailed; *domestic cow* alone contributes 2,109,009 images (approx. 54% of all samples). To reduce overfitting and speed up training, we cap each class at 100k images, following Tabak *et al.* (Tabak et al., 2020). This yields 816,495 balanced samples (Fig. 3) while preserving all 48 classes. The visualisation approach follows that of Perrett et al. (Perrett et al., 2023), where the dataset is divided into three groups: ‘Head’ classes (blue) contain > 50% of total samples, ‘Tail’ classes (pink) cover most species, while ‘Few-shot’ classes (green) include those with < 20 samples.

3.2 Data Preprocessing

Our workflow converts the raw image NACTI archive into an analysis-ready dataset in three stages.

(i) Automatic Animal Localisation. All images are first processed with MegaDetectorV6 to obtain bounding boxes and confidence scores. The resulting JSON files are merged with the original metadata such that each record includes both species labels and coordinates for the detected animal region.




Figure 4: **Training Dynamics.** Column (a) reports accuracy whilst column (b) reports on loss on one of the training runs with scheduling. Within each column, the left plot shows train and validation, right plots show the train-validation gap. Note the dashed lines marking early-stopping point.

(ii) **Class-Balanced Resampling.** Following Tabak *et al.*, we cap every category at 100,000 samples, discard empty/vehicle frames, and retain all 48 animal classes. This produces 816,495 images with a more balanced distribution profile shown in Figure 3.

(iii) **Split Strategy and Reduced-Bias Evaluation.** The balanced pool is split 80/10/10 (train/validation/test) with a fixed seed, acknowledging that visually similar burst information may still be spread across subsets in this setting. To gauge towards bias reduction and better generalisation we also construct and publish a *reduced-bias test set* using the ENA Dataset (LILA Science, 2024a), filtering to the nine classes shared with NACTI and removing any temporal/spatial overlap. All nine fall in the distribution ‘Tail’, providing a stringent domain-shift benchmark.

Final Pre-processing Steps. Images are resized to 256×256 px as per Tabak *et al.*, converted to tensors, and normalised with $\mu = [0.485, 0.456, 0.406]$, $\sigma = [0.229, 0.224, 0.225]$ to stabilise optimisation. Each training sample is finally stored as an information tuple (image, {bbox, label, conf}).

4 EXPERIMENTS

4.1 Baseline Model

Hardware Specification. All training is conducted utilising 11 NVIDIA Tesla P100 GPUs (Pascal architecture) to efficiently distribute computation across multiple devices and reduce training time. All inference was performed locally on an NVIDIA GeForce RTX 3090 GPU (Ampere architecture).

Full Network Training Details. All experiments fine-tune the ResNet-50 (He *et al.*, 2016) backbone of AI4GAmazonRainforest (Hernandez *et al.*, 2024) with its standard pre-trained weights, using AdamW (Loshchilov and Hutter, 2017) (weight-decay 10^{-2} , initial learning rate 10^{-4}) and a ReduceLROnPlateau scheduler that lowers the learning rate by factor 0.1 when validation metrics stagnate. To avoid premature termination or overfitting

Table 1: **NACTI Test Accuracy (LTR Methods vs. Baselines)** Comparison of (Overall/‘Tail’) accuracy on the biased NACTI test set across all configurations. Note that ‘Tail’ performance stats are saturated with the top accuracy of 99.22% is shared by three configurations.

LOSS + OPTIMISER CONFIGURATION	BASELINES (NO SCHEDULER)	LTR LR SCHEDULER
BASELINES		
Cross-Entropy + Adam	95.51% / 94.30%	Not Tested
Cross-Entropy + AdamW	96.77% / 95.64%	99.34% / 99.22%
LTR LOSSES		
Focal Loss + AdamW	96.65% / 95.96%	99.31% / 99.20%
WCE + AdamW	96.69% / 96.16%	99.35% / 99.22%
LDAM Loss + AdamW	96.71% / 96.17%	99.40% / 99.22%

caused by fluctuations in validation performance we employ early stopping (see Figure 4) with patience 10 and check-pointing, dynamically selecting the best model based on validation recall (see published code).

4.2 Experimental Results

Scheduler Impact on NACTI Results. Table 1 compares various combinations of LTR-sensitive learning rate scheduling and LTR loss functions+optimisers against baselines, evaluated in terms of both overall accuracy and performance on the ‘Tail’ classes. Introducing a learning rate scheduler from PyTorch (Paszke *et al.*, 2019) resulted in substantial performance gains across all configurations, with overall accuracy consistently exceeding 99% on the biased test set. The overall best result was achieved by combining LDAM loss with the AdamW optimiser and the ReduceLROnPlateau scheduler, yielding an accuracy of 99.40%—an improvement of approximately 2.66 percentage points over the same configuration without a scheduler (96.74%).

NACTI Tail Results. A similar trend was observed for the ‘Tail’ classes: accuracy increased from 96.17% to 99.22%, representing a relative improvement of 3.05%. Although all configurations achieved comparable performance on the ‘Tail’ classes following scheduler integration, the LDAM + AdamW com-

bination attained the highest overall accuracy and is henceforth referred to as *best (LTR) configuration*.

Reduced-Bias Test Set Results. All configurations are additionally evaluated on the independently constructed reduced-bias test set for assessing generalisation performance under distribution shift. Table 2 summarises this task showing expected, substantial accuracy drops relative to Table 1. While all configurations perform markedly worse in this OOD setting, using a learning rate scheduler yields consistent improvements across all loss functions. Among them, the LDAM loss combined with AdamW and a scheduler—our best configuration on the biased test set—achieves the highest accuracy (52.55%), demonstrating its robustness under both class imbalance and domain shift.

Table 2: **Reduced-Bias Test Dataset Accuracy.** Note that learning rate scheduling consistently improves performance and that the LDAM-Scheduler combination shows best performance in this OOD task.

LOSS + OPTIMISER CONFIGURATION	BASELINES (NO SCHEDULER)	LTR LR SCHEDULER
BASELINES		
Cross-Entropy with Adam	39.48%	Not Tested
Cross-Entropy with AdamW	38.27%	48.27%
LTR LOSSES		
Focal Loss with AdamW	37.62%	49.78%
WCE with AdamW	45.03%	51.20%
LDAM Loss with AdamW	38.77%	52.55%

Table 3: **Per-class Accuracy and Prevalence.** Results are shown for the Reduced-Bias Test Set under our best configuration.

CLASS	ACCURACY (%)	PREVALENCE
virginia opossum	65.79%	725
bobcat	61.26%	333
american black bear	77.05%	475
eastern gray squirrel	45.51%	312
wild turkey	0.00%	427
striped skunk	21.55%	297
red fox	54.72%	413
horse	0.00%	63
coyote	87.79%	344

5 DISCUSSION

5.1 Generalisation Considerations

Testing Bias and Performance Saturation. As discussed in Section 3.1, the standard NACTI Test Set – using the standard random selection process for testing NACTI so far – contains numerous visually similar frames originating from camera-trap burst capture, making it especially vulnerable to biased testing. As a result, test accuracy fails to reflect generalisability and leading to saturated 99%+ performances. However, the model’s ability to generalise is shown on the reduced-bias test set where different LTR loss




Figure 5: **Qualitative Example regarding Domain Shift for ‘Horse’.** Train-test contrast of some samples from the *horse* class where we observed catastrophic performance breakdown. Training images are often blurry or occluded, whereas reduced-bias test images involve night scenes, partial views and unusual poses.

functions can offer minor accuracy improvements as seen in Table 2. The dramatic performance gap between biased and reduced-bias test sets underscores the need for model-external OOD evaluation particularly for wildlife data like NACTI where significant context changes are part of environmental variability.

Per-Class Generalisation Analysis. To understand the generalisation behaviour of the best-performing model further on the Reduced-Bias Dataset, per-class accuracies for the 9 classes in this dataset are reported in Table 3. Note that all listed species belong to the ‘Tail’ of the training distribution. While the model demonstrates relatively strong performance on classes such as *coyote* and *american black bear*, accuracy varies considerably across species. Notably, some classes like *wild turkey* and *horse* receive 0% accuracy. These results underscore the difficulty of long-tailed generalisation under distribution shift. These findings suggest that while optimisation strategies such as LTR loss functions and learning rate schedules can improve overall accuracy, operation on rare classes may underperform catastrophically under distribution shift. Persistently low or zero accuracy on certain species across all configurations indicates that representational limitations play a more significant role unrecoverable by means of current learning alterations, potentially requiring a data-centric approach.

5.2 Failure Mode Analysis

Predictive Collapse for Tail Classes. Two classes – *horse* ($n = 63$) and *wild turkey* ($n = 148$) – consistently rank among the lowest-performing classes in the biased NACTI test set across *all* configurations and accuracy for both classes collapses to approximate 0% in the Reduced-Bias Test Set. Tables 4 exemplifies this collapse quantitatively. The model struggles to learn fine-grained distinctions between




Figure 6: **Qualitative Examples regarding Domain Shift for ‘Wild Turkey’.** Test images show poorer illumination, heavier occlusion and greater background clutter, widening the domain gap.

semantically or morphologically similar species when the training data does not contain sufficient examples of classes that actually contain species-determining features. This poses an ill-posed task of recognition of species who have little to no definition during the training phase in the first place. The consistent yet unsurprising misclassification into frequent categories confirms a strong prediction bias toward ‘Head’ classes, reflecting the model’s defaulting to dominating features learned from majority categories.

Qualitative Evidence. Figures 5 and 6 juxtapose representative images from the NACTI training set and the Reduced-Bias Test Set. Training samples are often blurry or partly occluded, whereas test images often contain only partial views or nocturnal infrared captures. Two interacting factors contribute to these systematic errors:

1. **Class-Imbalance Bias.** Misclassifications of severely under-represented classes in the training distribution occur due to network defaulting to dominant visual cues learned from the majority classes, yielding skewed posteriors that rely on class distribution priors rather than class-indicative image data.
2. **Domain Shift.** Compared with training images, our Reduced-Bias Test shots exhibit different lighting, more partial occlusion, night scenes and stronger motion blur overall. These discrepancies widen the representational gap and can lead to total detection collapse (see Table 4).

Partial View Considerations. For many tail classes including *horse*, the network almost invariably predicts dominant classes, for *horse* mainly *red deer* or *domestic cow* whenever a full body is not visible. This not only reflects the model’s bias towards the most frequent ‘Head’ classes under uncertainty, but points particularly to a problem of ‘unseen’ views. This observation points in the direction of future work to analyse data augmentation and the simulation

of partial view detection and related generative data strategies to compliment LTR scheduling and training methodologies.

6 CONCLUSIONS

Reproducible NACTI SOTA Performance. This paper delivers a reproducible improvement in long-tailed species recognition for the NACTI dataset. To make these capabilities accessible to the community, we publish the end-to-end pipeline that processes raw NACTI frames and fine-tunes a ResNet-50-based architecture with LTR losses under Adam/AdamW optimisers. We confirm experimentally that a scheduled LDAM/AdamW configuration will yield overall 99.4% Top-1 accuracy, setting a new reference point for NACTI wildlife species recognition under long-tailed conditions. We publish full code, weights, and complete data split information for reproducibility.

Table 4: **Collapse of Classification for ‘horse’.** Distribution of class recognitions (all false) of best performing LTR model in the Reduced-Bias Test Set confirms catastrophic collapse of classification with the classifier defaulting to head classes even under maximal use of LTR adjustments.

PREDICTED CLS	COUNT	PERCENT (%)
domestic cow	28	44.4
red deer	19	30.2
coyote	6	9.5
cougar	4	6.3
american black bear	3	4.8
unidentified deer	1	1.6
mule deer	1	1.6
unidentified bird	1	1.6

OOD Analysis and Recommendation. We also provide results on domain shifted classifier application for the ENA-DETECTION Dataset, highlighting both the generalisation improvements and the large remaining generalisation gap in such wildlife OOD experiments which current technology is not able to overcome without further training information. For all tested settings, we confirm consistent improvements of LTR-enhancing scheduling in and out of domain demonstrating best generalisation capabilities when scheduling is combined with LDAM losses. Thus, given these insights on NACTI in context with previous work (Alshammari et al., 2022) we recommend this setup as a basic starting point for addressing LTR tasks in wildlife monitoring. We hope that the publication of a reproducible SOTA pipeline for NACTI encourages further research in the evolving and important field of computer vision for wildlife monitoring.

REFERENCES

- Alshammari, S., Wang, Y.-X., Ramanan, D., and Kong, S. (2022). Long-tailed recognition via weight balancing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6897–6907.
- Beery, S., Morris, D., Yang, S., Simon, M., Norouzzadeh, A., and Joshi, N. (2019). Efficient pipeline for automating species id in new camera trap projects. *Biodiversity Information Science and Standards*, 3:e37222.
- Cao, K., Wei, C., Gaidon, A., Aréchiga, N., and Ma, T. (2019). Learning imbalanced datasets with label-distribution-aware margin loss. In *Neural Information Processing Systems*.
- Chen, G., Liaw, A., and Breiman, L. (2014). Class imbalance in ecology and conservation: A review of methods and applications. *Ecol. Informatics*, 24:35–43.
- Cui, Y., Jia, M., Lin, T.-Y., Song, Y., and Belongie, S. J. (2019). Class-balanced loss based on effective number of samples. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9260–9269.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., and Houlsby, N. (2020). An image is worth 16x16 words: Transformers for image recognition at scale. *ArXiv*, abs/2010.11929.
- Gadot, T., Istrate, S., Kim, H., Morris, D., Beery, S., Birch, T., and Ahumada, J. (2024). To crop or not to crop: Comparing whole-image and cropped classification on a large dataset of camera trap images. *IET Computer Vision*, 18(8):1193–1208.
- He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.
- Hernandez, A., Miao, Z., Vargas, L., Beery, S., Dodhia, R., and Lavista, J. (2024). Pytorch-wildlife: A collaborative deep learning framework for conservation.
- Jia, L., Tian, Y., and Zhang, J. (2022). Domain-aware neural architecture search for classifying animals in camera trap images. *Animals : an Open Access Journal from MDPI*, 12.
- Kang, B., Xie, S., Rohrbach, M., Yan, Z., Gordo, A., Feng, J., and Kalantidis, Y. (2020). Decoupling representation and classifier for long-tailed recognition. In *Eighth International Conference on Learning Representations (ICLR)*.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In Pereira, F., Burges, C., Bottou, L., and Weinberger, K., editors, *Advances in Neural Information Processing Systems*, volume 25. Curran Associates, Inc.
- Kühl, H. S. and Burghardt, T. (2013). Animal biometrics: quantifying and detecting phenotypic appearance. *Trends in ecology & evolution*, 28(7):432–441.
- LILA Science (2024a). ENA24-detection. <https://lila.science/datasets/ena24detection>. Accessed: 23 Oct 2025.
- LILA Science (2024b). NACTI Dataset: North American Camera Trap Images. <https://lila.science/datasets/nacti>. Accessed: 23 Oct 2025.
- Lin, T.-Y., Goyal, P., Girshick, R. B., He, K., and Dollár, P. (2017). Focal loss for dense object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42:318–327.
- Loshchilov, I. and Hutter, F. (2017). Decoupled weight decay regularization. In *International Conference on Learning Representations*.
- Norouzzadeh, M. S., Morris, D., Beery, S., Joshi, N., Jojic, N., and Clune, J. (2021). A deep active learning system for species identification and counting in camera trap images. *Methods in Ecology and Evolution*, 12(1):150–161.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. (2019). Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035.
- Perrett, T., Sinha, S., Burghardt, T., Mirmehdi, M., and Damen, D. (2023). Use your head: Improving long-tail video recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2415–2425.
- 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 in ecology & evolution*, 40(2):191–207.
- Tabak, M. A., Norouzzadeh, M. S., Wolfson, D. W., Newton, E. J., Boughton, R. K., Ivan, J. S., Odell, E. A., Newkirk, E. S., Conrey, R. Y., Stenglein, J. L., Ianarilli, F., Erb, J., Brook, R. K., Davis, A. J., Lewis, J. S., Walsh, D. P., Beasley, J. C., VerCauteren, K. C., Clune, J., and Miller, R. S. (2020). Improving the accessibility and transferability of machine learning algorithms for identification of animals in camera trap images: mlwic2. *Ecology and Evolution*, 10(22):12757–12767.
- Tabak, M. A., Norouzzadeh, M. S., Wolfson, D. W., Sweeney, S. J., VerCauteren, K. C., Snow, N. P., Halseth, J. M., Di Salvo, P. A., Lewis, J. S., Teton, A. M., Boughton, R. K., and Miller, R. S. (2022). Cameratrappdetector: Automatically detect, classify, and count animals in camera trap images using artificial intelligence. *bioRxiv*.
- Tuia, D., Kellenberger, B., Beery, S., Costelloe, B. R., Zuffi, S., Risse, B., Mathis, A., Mathis, M. W., Van Langevelde, F., Burghardt, T., et al. (2022). Perspectives in machine learning for wildlife conservation. *Nature communications*, 13(1):1–15.
- Zhang, Y., Kang, B., Hooi, B., Yan, S., and Feng, J. (2021). Deep long-tailed learning: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45:10795–10816.