

Overview of AnimalCLEF 2026: Discovery and Re-identification of Individual Animals

Lukáš Adam¹, Kostas Papafitsoros², Dean A. Williams³, Daniella Biffi³ and Lukas Pícek^{1,4}

¹University of West Bohemia in Pilsen, Pilsen, Czech Republic

²Queen Mary University of London, London, United Kingdom

³Texas Christian University, Fort Worth, Texas, USA

⁴Massachusetts Institute of Technology, Cambridge, USA

Abstract

The second edition of the individual animal identification challenge, AnimalCLEF 2026, organized within LifeCLEF, advances the field of animal re-identification using computer vision and machine learning. Incorporating new data into the previous year's edition, AnimalCLEF 2026 challenges participants to group images of unknown individuals so that each cluster corresponds to a single animal. The mix of four species, lynxes, salamanders, sea turtles and lizards, with different image capture conditions, attempts to make the submitted prediction models generalizable to unseen species. The competition attracted 283 participants across 230 teams, with 137 teams outperforming the provided baseline using the combination of MegaDescriptor and MiewID. This overview paper provides (i) a comprehensive description of the challenge and provided baseline method, (ii) detailed characteristics of the dataset and task specifications, (iii) an examination of the methods employed by contestants, and (iv) a discussion of the competition outcomes. The results highlight incremental advancements in animal re-identification, showcasing innovative approaches and techniques that push the limits of previous work.

Keywords

LifeCLEF, AnimalCLEF, fine-grained visual categorization, metadata, discovery, animals, re-identification, individual identification, machine learning, computer vision, classification

1. Introduction

Animal re-identification and individual discovery are becoming increasingly important tools for wildlife monitoring and ecological research. While traditional re-identification focuses on matching observations of known individuals, clustering-based individual discovery aims to automatically group images belonging to the same individual animal without requiring predefined identity labels. These tasks support a wide range of applications, including disease monitoring and control [1], estimating population abundance [2], tracking invasive species [3], and assessing the impact of human activities on ecosystems and restoration efforts [4]. Advances in camera trapping and citizen science have led to rapidly growing wildlife image collections, increasing the need for scalable automated approaches capable of analyzing large amounts of visual data [5, 6]. Automating the re-identification and discovery process enables large-scale analysis of population dynamics, migration routes, habitat utilization, and animal behavior. Such information can support survival estimation [7], behavioral studies [8], injury and disease monitoring [9], measuring ecotouristic pressure [10], or estimating species distribution [11]. Compared to supervised re-identification, clustering-based approaches are particularly attractive because they reduce the dependence on manually annotated identity labels, which are often difficult to obtain in ecological settings. Despite substantial progress in machine learning and computer vision, discovering animal individuals in realistic ecological settings remains a difficult problem due to several factors. Some individuals are represented by hundreds of observations, while others may appear only once or twice. Moreover, different individuals can appear more visually similar to each other than a single individual does to itself across changing poses, viewpoints, or environmental conditions. These

CLEF 2026 Working Notes, 21 – 24 September 2026, Jena, Germany

*Corresponding author.

✉ adamluk3@fel.zcu.cz (L. Adam); picekl@kky.zcu.cz (L. Pícek)

ORCID 0000-0001-8748-4308 (L. Adam); 0000-0001-9691-4576 (K. Papafitsoros); 0000-0002-9001-6019 (D. A. Williams); 0000-0001-6437-9855 (D. Biffi); 0000-0002-6041-9722 (L. Pícek)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).




Figure 1: Individual animals similarities. Visual similarities and differences across three individuals (columns) from each of the three datasets in AnimalCLEF 2026.

challenges make clustering particularly difficult, as methods must simultaneously learn discriminative representations and infer the underlying population structure without supervision. Developing benchmarks and evaluation protocols for this task is therefore important not only for measuring algorithmic progress but also for encouraging methods that are robust, scalable, and applicable across diverse species and monitoring scenarios.

To address these issues and enable continual incremental improvements in animal re-identification, we organize the AnimalCLEF research competition. The first 2025 edition focused on the open-set identification of individual animals. The 2026 edition, which took part in the LifeCLEF 2026 [12] and the FGVC13 workshop at CVPR 2026, builds directly upon the foundations established in the 2025 challenge and extends the problem setting from open-set identification to unsupervised individual discovery and clustering. Instead of matching query images to a gallery of known identities, participants must automatically group images depicting the same individual without prior knowledge of the number or identities of animals in the dataset. This setting better reflects realistic ecological scenarios, where identity annotations are often unavailable and new individuals continuously appear in the data. The 2026 dataset also expands the species coverage. In addition to lynxes, salamanders, and sea turtles, it introduces Texas horned lizards as a new taxonomic group (see Figure 1).

To provide participants with a starting point, we provided a baseline at the start of the challenge and an improved version midway through. The initial baseline was based on the deep models MegaDescriptor [13] and MiewID [14] with DBSCAN [15], while the second baseline combined them by WildFusion [16] and HDBSCAN [17]. We consider this competition a big success, as both baselines were outperformed by 132 teams, bringing new ideas to the field of animal re-identification.

2. Challenge Description

2.1. Dataset

The AnimalCLEF 2026 dataset builds upon the previous editions of LifeCLEF [18, 19], AnimalCLEF [20], and the WildlifeReID-10k dataset [21]. The WildlifeReID-10k dataset was created by using the Python package wildlife-datasets [13] by combining the following datasets [5, 22–50]. All of the images are cropped to contain the part used to recognize individual animals (heads for sea turtles, bodies for tigers, tails for whales, . . .). Some of the images were supported by additional metadata, including timestamp, orientation, and clusters of similar images.

The AnimalCLEF 2026 dataset is a direct extension of the AnimalCLEF 2025 dataset [20], which included Eurasian lynxes *Lynx lynx*, loggerhead sea turtles *Caretta caretta*, and fire salamanders *Salamandra salamandra*. The lynx data come from the CzechLynx dataset [51], while the sea turtle data originate from an unreleased subset of the SeaTurtleID2022 dataset [23]. The salamander data have not yet been publicly released. For 2026, the dataset has been extended with Texas horned lizards *Phrynosoma cornutum*, whose images were previously released by [52], but without identity labels.

The subset with **Eurasian lynxes** (*Lynx lynx*) individuals includes camera-trap photographs, with annotations for three computer-vision tasks: individual identification, pose estimation, and instance segmentation [51]. It spans more than 15 years, covers two regions in Central Europe (Southwest Bohemia and the Western Carpathians), and contains around 37k manually verified and labeled images.

The subset with **fire salamander** (*Salamandra salamandra*) individuals consists of images collected under specific environmental conditions. As these amphibians are primarily active at night and only under high humidity, most captures occurred during rainy nights. Within a fixed area, salamanders were located using visual encounter surveys conducted by a single, consistent observer throughout the study. Each individual was then photographed from dorsal, ventral, and right lateral perspectives. Following standard protocol, all salamanders were released immediately at the site of capture.

The subset with **loggerhead sea turtles** (*Caretta caretta*) individuals comprises 8,729 photographs of 438 individuals collected over 13 years on Zakynthos Island, Greece. Each photograph includes various annotations, e.g., identity, encounter timestamp, and body parts segmentation masks. As the database, we considered the 8,729 images from the original publication [23]. As a query set, we also published 500 images of 100 individuals: 50 were already in the database, and 50 were new.

The subset with **Texas horned lizards** (*Phrynosoma cornutum*) comes from a long-term study conducted from 2015 to 2021 in two towns in south Texas [52]. In addition to collecting photos of lizards with clearly visible ventral spots, the authors also collected genotypes for individual animals and confirmed that the patterns remained stable over time.

The split into the training and test sets follows the previous edition (where it was called database and query due to a different task). Due to its small size, all lizard images were put in the test set. WildlifeReID-10k was proposed as an additional training set. The four datasets form a balanced mix of image-to-identity ratios. The majority of identities in the salamander and lizard datasets are present in only one image. The lynx dataset has a more balanced image-to-identity ratio, whereas the sea turtle dataset contains 100 identities, each with 5 images. Detailed information is provided in Table 1.

Table 1

AnimalCLEF 2026 dataset statistics. We list the number of images and individuals in each subset. The *known* and *unknown* represent how many individuals were and were not available in the training set.

Subset	Number of Images	Number of Individuals		
		Total	Known	Unknown
WildlifeReID-10k (<i>Additional Training</i>)	140,488	10,772	10,772	-
<i>Lynx lynx</i>	3,903	99	99	-
↳ <i>Training</i>	2,957	77	77	-
↳ <i>Test</i>	946	41	19	22
<i>Salamandra salamandra</i>	2,077	646	646	-
↳ <i>Training</i>	1,388	587	587	-
↳ <i>Test</i>	689	538	479	59
<i>Caretta caretta</i>	9,229	488	488	-
↳ <i>Training</i>	8,729	438	438	-
↳ <i>Test</i>	500	100	50	50
<i>Phrynosoma cornutum</i>	274	217	-	-
↳ <i>Test</i>	274	217	0	217

The test set was further split into public and private subsets. Whenever a participant made a submission, they received an evaluation on the public test set. The final evaluation was performed on the private test set. The participant did not know the split into private and public test sets.

2.2. Task and Evaluation Protocol

The competitors were required to divide the test set into clusters, with each cluster corresponding to a single identity. The provided metadata listed the animal species; the clustering was performed dataset-wise. The evaluation metric was the adjusted Rand index (ARI). Considering all possible pairs of samples, let true positives (tp) denote the number of pairs predicted to belong to the same cluster that indeed belong to the same cluster (are the same identity), true negatives (tn) the number of pairs predicted to belong to different clusters that indeed belong to different clusters (are different identities), and similarly false positives (fp) and false negatives (fn). The Rand index (RI) is defined by

$$\text{RI} = \frac{\text{tp} + \text{tn}}{\text{tp} + \text{tn} + \text{fp} + \text{fn}}.$$

Since this number can be close to 1 even for trivial clusterings (such as each cluster having one image), we employed the adjusted Rand index (ARI) defined by

$$\text{ARI} = \frac{2(\text{tp} \cdot \text{tn} - \text{fn} \cdot \text{fp})}{(\text{tp} + \text{fn})(\text{fn} + \text{tn}) + (\text{tp} + \text{fp})(\text{fp} + \text{tn})}.$$

This metric takes values in $[-1, 1]$, where 1 corresponds to a perfect clustering, and 0 corresponds to a random assignment. Moreover, it is invariant under permutations of cluster labels and identity ordering. To properly balance the datasets, we computed ARI for each dataset separately, and the final score was defined as the mean ARI.

2.3. Baselines

To enable a smoother start for all participants and straightforward model evaluation, we provided two baselines. Both were based on the animal re-identification models [MegaDescriptor-L-384](#) [13] and [MiewID](#) [14]. The first baseline assigned one of these models to each dataset. The models extracted a feature vector for each test image. The clustering was performed on these extracted features by DBSCAN [15] with a manually selected threshold ε . To account for the specificity of animal datasets, we decreased the minimum number of images per cluster to two. Since this procedure required the manual tuning of the assigned model and the ε threshold, during the competition, we provided another baseline with [WildFusion](#) [16] for an automatic selection of the model and [HDBSCAN](#) [17] for an automatic tuning of the ε hyperparameter.

Both models achieved relatively strong performance in animal re-identification, serving as a robust starting point. Overall, our goal was to provide a comprehensive and accessible starting point for all participants, enabling them to focus on developing novel solutions and improving baselines.

2.4. Timeline

The AnimalCLEF 2026 competition was launched on February 4, 2026, and was promoted through the [LifeCLEF](#) [12], [Kaggle](#), and [FGVC](#) challenge web pages, inviting participants to register. The competition ran for approximately three months, with the final submission deadline on April 30. Similar to other LifeCLEF challenges, the test data remained confidential. Participants were allowed to make up to five submissions per day using the Kaggle evaluation platform to assess their models on the public query set. Based on these results, the participants were allowed to select five submissions, which were evaluated on the private test set. The participants were scored and ranked based on the better of the two selected solutions.

2.5. Working Notes

Participants were encouraged to submit both their code and a detailed technical report (Working Notes) to ensure their results can be fully reproduced. All submitted Working Notes underwent thorough review and received detailed feedback from 2-3 experts with a strong background in CV and ML. This rigorous review process was designed to guarantee reproducibility and maintain quality standards. The review was single-blind, allowing participants to respond with up to two rebuttals to address any feedback or concerns raised by the reviewers. These working notes provide an in-depth analysis of the techniques employed, including hyperparameter tuning, model ensembling, and loss function selection, offering valuable insights into the development of the method for animal re-identification.

3. Challenge Results

AnimalCLEF 2026 attracted 283 participants across 230 teams, a slightly higher number than the previous year. 137 teams outperformed the MegaDescriptor + MiewID + DBSCAN baseline (with ARI 0.206), and 132 teams outperformed the WildFusion + HDBSCAN baseline (with ARI 0.212).

While the official final scoring was based on the submissions selected by each team, as described in Section 2.4, our analysis uses all submitted runs. This allows us to provide a more detailed view of participant performance, but may lead to slight differences from the official leaderboard. The team rankings are taken from the official ranking.

Figure 2 shows the results of the best 50 participants on the private test set. The first 4 graphs show performance on the individual datasets for each animal, and the last one shows their average. The blue lines show the ARIs for all teams (lower ranks are better), and the orange line shows their averages. The last graph should show a decreasing function; however, there are a few kinks due to the final submission selection process described above. The ARI on sea turtles is relatively stable for all teams. The ranking of teams was most influenced by their performance on salamanders and lizards. Since these are the two datasets with the lowest image-to-identity ratios, the ARI metric appears highly sensitive. Three teams (M.I.A, VIPL-VSU and HSU Lab) solved the lizard dataset perfectly in at least one submission.




Figure 2: Private Leaderboard of the AnimalCLEF 2026 competition. We report the performance of the top 50 teams (out of 230). The track metric is ARI (adjusted Rand index) on individual datasets. The orange line is the mean performance of the 50 teams.

Figure 3 shows the population estimates (number of clusters) of the best 50 participants on the whole test set. The solid orange line shows the true population size. The population estimates are relatively precise for salamanders and lizards, but overshoot significantly for lynxes and sea turtles. A probable explanation is that the images for the former two datasets were captured from a single or a few angles, whereas those for the latter two were captured from multiple angles. Even though there are some similarities between opposite profiles for sea turtles [22], this does not seem to hold for carvinorees [53]. In other words, merging opposite profiles cannot be reliably done, and the estimated population size may be twice the true population size. Indeed, we see that the population estimates follow much better the doubled true population size (dashed orange line). We also numerically summarise the population estimates for the top 5 performing teams in Table 2.




Figure 3: Population estimates for the AnimalCLEF 2026 competition. We report population estimates (number of clusters) for the top 50 teams (out of 230). The solid orange line is the true population size. Since lynxes and sea turtles have different patterns on opposite sides, the population estimates may be twice the true population sizes. We show double the true population size by the dashed orange line.

Table 2

Population estimates for 5 top-performing teams. Each row corresponds to one dataset, each column to one team, with the last row corresponding to the true population size.

	lynxes	salamanders	sea turtles	lizards
M.I.A	138	594	179	204
Ram Lab	80	559	176	205
VIPL-VSU	21	551	153	213
HSU Lab	218	452	167	213
DS@GT LifeCLEF	216	574	219	220
TRUE	41	538	100	217

Table 3 shows a more detailed analysis of the top 5 performing teams, which were all within 3% of ARI. The coloured blocks show the confusion tables for each team on each dataset. The top-left and bottom-right are true negatives (tn) and true positives (tp), while the top-right and bottom-left are false positives (fp) and false negatives (fn); see their definitions in Section 2.2. The goal was to maximize the diagonal elements and minimize the off-diagonal elements. Following the discussion in [54], these off-diagonal elements have a different significance. False positives, also called “false match”, merge two identities into a single cluster. False negatives, also called “missing link”, split one identity into two clusters. False matches are usually more difficult to detect and fix, as they require manually analyzing the entire cluster to determine which image corresponds to each merged identity. Interestingly, DS@GT LifeCLEF produced only a small number of false matches, indicating that their clusters were very pure.

Table 3

Confusion tables for 5 top-performing teams. Each 2×2 confusion table shows the true negatives (top left), true positives (bottom right), false positives (top right), also known as false matches, and false negatives (bottom left), also known as missing links.

	M.I.A		Ram Lab		VIPL-VSU		HSU Lab		DS@GT LifeCLEF	
lynxes	817410	4020	819328	2102	736998	84432	819020	2410	821316	114
	27750	44790	35044	37496	25856	46684	43308	29232	42018	30522
salamanders	473034	26	472980	80	472808	252	472502	558	472996	64
	626	346	356	616	554	418	534	438	516	456
sea turtles	247144	356	247054	446	246856	644	247260	240	247486	14
	1044	956	990	1010	748	1252	626	1374	944	1056
lizards	74570	54	74586	38	74598	26	74602	22	74604	20
	2	176	0	178	0	178	0	178	18	160

4. Participants and Methods

We have received 15 working notes submissions, of which 1 was desk-rejected due to a late submission date. The remaining received 2-3 reviews, and most authors opted in for an additional optional AI review. 12 working notes were accepted into the proceedings. Below, we provide a short summary of the working notes of participants within the top 50 teams. Details of all submitted methods and systems are further developed in participants’ working notes [55–66].

Team **M. I. A** [55] (Top 1) used species-specific preprocessing by segmenting the animals and highlighting the important body parts. Global similarities were computed using an ensemble of MiewID models, including variants fine-tuned with Dynamic ArcFace and SphereFace2-Focal losses, and combined with local feature matching from ALIKED+LightGlue and DISK+LightGlue using calibrated score fusion (WildFusion). Query images were clustered using K-Reciprocal Re-Ranking and Chinese Whispers, and confident matches were converted into clusters. The hyperparameters were selected using Optuna.

Team **Ram Lab** [56] (Top 2) built independent pipelines for each species instead of a common model. Depending on the species, the method combined fine-tuned global embeddings, local feature matching, and pretrained descriptors, followed by hierarchical agglomerative clustering with species-specific thresholds or target cluster counts. Salamanders relied mainly on segmented local pattern matching supplemented by a self-trained embedding model to bridge viewpoint differences, while Texas horned lizards used specialized methods suited to ventral patterns.

Team **HSU Lab** [57] (Top 4) applied a two-phase pipeline. The first phase was general across all species and used HDBSCAN to cluster the similarity scores obtained by MiewID, MegaDescriptor, and SIFT+LightGlue. The second phase refined the initial clusters by species-specific correction stages, including orientation-aware merging, belly-region graph replacement, region-normalized merging, or signed relation repair.

Team **DS@GT LifeCLEF** [58] (Top 5) applied segmentation with species-specific padding and extracted global features to retrieve candidate pairs for local features SuperPoint, SIFT, ALIKED and DISK. The obtained high-similarity pairs were verified by a trained LightGBM binary classifier to remove high-scoring false pairs that wrongly merge distinct individuals. Only high-confidence edges were retained before applying Leiden community detection for clustering.

Team **gilles colling** [59] (Top 6) combined MegaDescriptor, MiewID, and LightGlue local matching to compute pairwise similarities. Test images sufficiently similar to reference identities were anchored to known identities, while non-anchored image pairs were merged only when a pairwise XGBoost

classifier predicted a high-confidence match. The pair classifiers were trained on manually annotated same/different test–test and test–train image pairs, and additionally hand-labeled spot keypoints for salamanders. Final identity clusters were obtained as connected components of the graph.

Team **ymygan** [60] (Top 33) used MegaDescriptor and MiewID embeddings to extract similarity matrices. The clustering was species-specific and combined reciprocal-neighbor filtering, pseudo-label refinement, partition co-association, descriptor ensembling, and initialization ensembling. Texas horned lizards were processed using density-based clustering because no reference identities were available.

Team **Hanif Noer Rofiq** [61] (Top 43) proposed Lineage-Aware Markov Partition Refinement (LAMPR), which started from an initial clustering based on global descriptors and refined it through iterative partition edits. Local feature matching, geometric verification, size priors, and orientation-normalized matching were used selectively to modify only uncertain regions of the partition. Different refinement rules were applied to each species based on its characteristic identity cues.

Team **Dr_F** [62] (Top 49) extracted features from multiple large foundation models, including DINOv3, InternViT, SigLIP2, EVA02-CLIP, and MegaDescriptor, then concatenated their embeddings. Supervised contrastive projection heads were trained on the combined features, and several clustering algorithms were evaluated on the learned embedding space. The authors also experimented with ArcFace metric learning, pseudo-labeling, or lookup-based approaches.

5. Conclusions and Discussion

This paper presents an overview and evaluation of the second edition of the AnimalCLEF challenge, organized jointly with the CLEF LifeCLEF lab and the CVPR-FGVC13 workshop (The 13th Workshop on Fine-Grained Visual Categorization) held at CVPR. The challenge extended the three species from AnimalCLEF 2025—*Caretta caretta*, *Salamandra salamandra*, and *Lynx lynx*—by introducing a fourth species, *Phrynosoma cornutum*. While the 2025 edition required participants to determine whether a query image belonged to a known individual or represented a previously unseen individual, the 2026 edition further required participants to group images of unseen individuals into identity clusters. The two tasks are complementary: the former distinguishes known from unknown individuals, whereas the latter identifies the unknown individuals themselves. The competition attracted 283 participants (compared to 207 in 2025) and 15 working notes (compared to 8 in 2025), describing a diverse range of methods and ideas. The main observations from this edition are summarized below:

- **Global-local feature fusion remained the dominant paradigm.** Nearly all methods combined global embeddings (typically MegaDescriptor and MiewID) with local feature matching (ALIKED, DISK, SIFT, or SuperPoint), using either WildFusion or manually designed score fusion.
- **Species-specific pipelines became standard practice.** Rather than using a single model for all datasets, most top-performing methods employed species-specific preprocessing, thresholds, and post-processing tailored to the visual characteristics of each species.
- **Graph-based clustering replaced direct identity prediction.** Most successful approaches formulated the task as graph construction, where pairwise similarities were converted into graph edges and clustered using algorithms such as Chinese Whispers, Leiden, HDBSCAN, connected components, or hierarchical agglomerative clustering.
- **Pair verification became an important refinement step.** Top methods trained dedicated pairwise classifiers (e.g., XGBoost, LightGBM, or Random Forest) that combined global descriptors, local matching scores, and handcrafted features to filter uncertain image pairs before clustering.
- **Generalization to ecological metrics remained challenging.** Although the best methods achieved high ARI scores, they were substantially less accurate at estimating biologically relevant quantities such as the number of distinct individuals, indicating that improvements in clustering quality do not necessarily translate into improved ecological population estimates.

AnimalCLEF continues to evolve alongside the rapidly advancing field of animal re-identification. Many of the ideas introduced by participants in this edition, such as graph-based clustering, species-specific processing, and learned pairwise verification, are likely to influence future research beyond the competition itself. As organizers, we strongly encouraged participants to release their source code and trained models, allowing the community to reproduce, compare, and build upon the proposed methods. We expect that these openly available implementations will accelerate the transfer of competition solutions into practical wildlife monitoring systems and support future research on large-scale, non-invasive biodiversity monitoring. Future editions of AnimalCLEF will continue to expand the benchmark with new species, more challenging scenarios, and evaluation protocols that better reflect real-world ecological applications.

Declaration on Generative AI

During the preparation of this work, the authors used Grammarly for grammar and spelling checks and ChatGPT for improving clarity and rewording sentences. After using this tool/service, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

Acknowledgments

LP was supported by the Technology Agency of the Czech Republic, project No. SS73020004. LA was supported by the Center for advanced machines and manufacturing technology (CAMAT) - TN02000028. We would like to thank all participants for sharing their codebases and working nodes, which are available at <https://www.kaggle.com/c/animal-clef-2026>.

References

- [1] P. Palencia, R. Vada, S. Zanet, M. Calvini, A. De Giovanni, G. Gola, E. Ferroglio, Not just pictures: Utility of camera trapping in the context of african swine fever and wild boar management, *Transboundary and Emerging Diseases* 2023 (2023) 1–9.
- [2] J. M. Rowcliffe, J. Field, S. T. Turvey, C. Carbone, Estimating animal density using camera traps without the need for individual recognition, *Journal of Applied Ecology* (2008) 1228–1236.
- [3] A. Caravaggi, M. Zaccaroni, F. Riga, S. C. Schai-Braun, J. T. Dick, W. I. Montgomery, N. Reid, An invasive-native mammalian species replacement process captured by camera trap survey random encounter models, *Remote Sensing in Ecology and Conservation* 2 (2016) 45–58.
- [4] J. D. Blount, M. W. Chynoweth, A. M. Green, Ç. H. Şekercioğlu, COVID-19 highlights the importance of camera traps for wildlife conservation research and management, *Biological Conservation* 256 (2021) 108984.
- [5] J. R. Parham, J. Crall, C. Stewart, T. Berger-Wolf, D. Rubenstein, Animal population censusing at scale with citizen science and photographic identification, in: 2017 AAAI Spring Symposium Series, 2017.
- [6] S. Schneider, G. W. Taylor, S. Linquist, S. C. Kremer, Past, present and future approaches using computer vision for animal re-identification from camera trap data, *Methods in Ecology and Evolution* 10 (2019) 461–470.
- [7] G. Schofield, M. Klaassen, K. Papafitsoros, M. Lilley, K. A. Katselidis, G. C. Hays, Long-term photo-id and satellite tracking reveal sex-biased survival linked to movements in an endangered species, *Ecology* 11 (2020) e03027.
- [8] G. Schofield, K. Papafitsoros, C. Chapman, A. Shah, L. Westover, L. C. D. Dickson, K. A. Katselidis, More aggressive sea turtles win fights over foraging resources independent of body size and years of presence, *Animal Behaviour* 190 (2022) 209–219.
- [9] K. Papafitsoros, A. Panagopoulou, G. Schofield, Social media reveals consistently disproportionate tourism pressure on a threatened marine vertebrate, *Animal Conservation* 24 (2021) 568–579.

- [10] K. Papafitsoros, L. Adam, G. Schofield, A social media-based framework for quantifying temporal changes to wildlife viewing intensity, *Ecological Modelling* 476 (2023) 110223.
- [11] M. E. Hanna, E. M. Chandler, B. X. Semmens, T. Eguchi, G. E. Lemons, J. A. Seminoff, Citizen-sourced sightings and underwater photography reveal novel insights about green sea turtle distribution and ecology in southern california, *Frontiers in Marine Science* 8 (2021) 671061.
- [12] L. Pícek, L. Adam, S. Kahl, R. Bossy, L. Chrobak, H. Goëau, K. Papafitsoros, H. Klinck, W.-P. Vellinga, R. Planqué, T. Denton, K. Barnard, C. Nédellec, L. Deléger, M. Courtin, G. Martellucci, I. Moummad, F. Vinatier, P. Bonnet, A. Joly, Overview of LifeCLEF 2026: Ai challenges for biodiversity understanding and ecosystem management, in: *International Conference of the Cross-Language Evaluation Forum for European Languages (CLEF)*, Springer, 2026.
- [13] V. Čermák, L. Pícek, L. Adam, K. Papafitsoros, WildlifeDatasets: An Open-Source Toolkit for Animal Re-Identification, in: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2024, pp. 5953–5963.
- [14] L. Otarashvili, T. Subramanian, J. Holmberg, J. Levenson, C. V. Stewart, Multispecies animal re-id using a large community-curated dataset, *arXiv preprint arXiv:2412.05602* (2024).
- [15] M. Ester, H.-P. Kriegel, J. Sander, X. Xu, et al., A density-based algorithm for discovering clusters in large spatial databases with noise, in: *kdd*, volume 96, 1996, pp. 226–231.
- [16] V. Cermak, L. Pícek, L. Adam, L. Neumann, J. Matas, Wildfusion: Individual animal identification with calibrated similarity fusion, in: *European Conference on Computer Vision*, Springer, 2025, pp. 18–36.
- [17] L. McInnes, J. Healy, Accelerated hierarchical density based clustering, in: *2017 IEEE international conference on data mining workshops (ICDMW)*, IEEE, 2017, pp. 33–42.
- [18] A. Joly, H. Goëau, S. Kahl, L. Pícek, C. Botella, D. Marcos, M. Šulc, M. Hruz, T. Lorieul, S. S. Moussi, et al., Lifeclef 2023 teaser: Species identification and prediction challenges, in: *European Conference on Information Retrieval*, Springer, 2023, pp. 568–576.
- [19] L. Pícek, S. Kahl, H. Goëau, L. Adam, T. Larcher, C. Leblanc, M. Servajean, K. Janoušková, J. Matas, V. Čermák, K. Papafitsoros, R. Planqué, W.-P. Vellinga, H. Klinck, T. Denton, J. S. Cañas, G. Martellucci, F. Vinatier, P. Bonnet, A. Joly, Overview of LifeCLEF 2025: Challenges on species presence prediction and identification, and individual animal identification, in: *International Conference of the Cross-Language Evaluation Forum for European Languages (CLEF)*, Springer, 2025.
- [20] L. Adam, K. Papafitsoros, R. Kovář, V. Čermák, L. Pícek, Overview of animalclef 2025: recognizing individual animals in images, *Working Notes of CLEF* (2025).
- [21] L. Adam, V. Cermak, K. Papafitsoros, L. Pícek, Wildlifereid-10k: Wildlife re-identification dataset with 10k individual animals, in: *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR) Workshops*, 2025, pp. 2099–2109.
- [22] L. Adam, K. Papafitsoros, C. Jean, A. F. Rees, V. Čermák, Exploiting facial side similarities to improve ai-driven sea turtle photo-identification systems, *Ecological Informatics* 89 (2025) 103158.
- [23] L. Adam, V. Čermák, K. Papafitsoros, L. Pícek, Seaturtleid2022: A long-span dataset for reliable sea turtle re-identification, in: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2024, pp. 7146–7156.
- [24] W. Andrew, J. Gao, S. Mullan, N. Campbell, A. W. Dowsey, T. Burghardt, Visual identification of individual holstein-friesian cattle via deep metric learning, *Computers and Electronics in Agriculture* 185 (2021) 106133.
- [25] W. Andrew, C. Greatwood, T. Burghardt, Visual localisation and individual identification of holstein friesian cattle via deep learning, in: *Proceedings of the IEEE International Conference on Computer Vision Workshops*, 2017, pp. 2850–2859.
- [26] W. Andrew, S. Hannuna, N. Campbell, T. Burghardt, Automatic individual holstein friesian cattle identification via selective local coat pattern matching in RGB-D imagery, in: *2016 IEEE International Conference on Image Processing (ICIP)*, IEEE, 2016, pp. 484–488.
- [27] J. Bruslund Haurum, A. Karpova, M. Pedersen, S. Hein Bengtson, T. B. Moeslund, Re-identification of zebrafish using metric learning, in: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision Workshops*, 2020, pp. 1–11.

- [28] Local Ocean Conservation, Turtle recall: Conservation challenge, 2021. URL: <https://zindi.africa/competitions/turtle-recall-conservation-challenge>.
- [29] N. Dlamini, T. L. van Zyl, Automated identification of individuals in wildlife population using siamese neural networks, in: 2020 7th International Conference on Soft Computing & Machine Intelligence (ISCMi), IEEE, 2020, pp. 224–228.
- [30] A. C. Ferreira, L. R. Silva, F. Renna, H. B. Brandl, J. P. Renoult, D. R. Farine, R. Covas, C. Doutrelant, Deep learning-based methods for individual recognition in small birds, *Methods in Ecology and Evolution* 11 (2020) 1072–1085.
- [31] A. Freytag, E. Rodner, M. Simon, A. Loos, H. S. Kühl, J. Denzler, Chimpanzee faces in the wild: Log-Euclidean CNNs for predicting identities and attributes of primates, in: German Conference on Pattern Recognition, Springer, 2016, pp. 51–63.
- [32] L. Fu, G. He, Cow dataset, 2021. URL: <https://doi.org/10.6084/m9.figshare.16879780.v1>.
- [33] J. Gao, T. Burghardt, W. Andrew, A. W. Dowsey, N. W. Campbell, Towards self-supervision for video identification of individual holstein-friesian cattle: The Cows2021 dataset, arXiv preprint arXiv:2105.01938 (2021).
- [34] Z. He, J. Qian, D. Yan, C. Wang, Y. Xin, Animal re-identification algorithm for posture diversity, in: ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE, 2023, pp. 1–5.
- [35] J. Holmberg, B. Norman, Z. Arzoumanian, Estimating population size, structure, and residency time for whale sharks *Rhincodon typus* through collaborative photo-identification, *Endangered Species Research* 7 (2009) 39–53.
- [36] D. Kern, T. Schiele, U. Klauck, W. Ingabire, Towards automated chicken monitoring: Dataset and machine learning methods for visual, noninvasive reidentification, *Animals* 15 (2024) 1.
- [37] M. Lahiri, C. Tantipathananandh, R. Warungu, D. I. Rubenstein, T. Y. Berger-Wolf, Biometric animal databases from field photographs: identification of individual zebra in the wild, in: Proceedings of the 1st ACM international conference on multimedia retrieval, 2011, pp. 1–8.
- [38] S. Li, J. Li, H. Tang, R. Qian, W. Lin, ATRW: A Benchmark for Amur Tiger Re-identification in the Wild, in: Proceedings of the 28th ACM International Conference on Multimedia, Association for Computing Machinery, 2020, p. 2590–2598.
- [39] T.-Y. Lin, Cat individual images, 2020. URL: <https://www.kaggle.com/datasets/timost1234/cat-individuals>.
- [40] V. Miele, G. Dussert, B. Spataro, S. Chamailé-Jammes, D. Allainé, C. Bonenfant, Revisiting animal photo-identification using deep metric learning and network analysis, *Methods in Ecology and Evolution* 12 (2021) 863–873.
- [41] G. Mougeot, D. Li, S. Jia, A deep learning approach for dog face verification and recognition, in: PRICAI 2019: Trends in Artificial Intelligence: 16th Pacific Rim International Conference on Artificial Intelligence, Cuvu, Yanuca Island, Fiji, August 26–30, 2019, Proceedings, Part III 16, Springer, 2019, pp. 418–430.
- [42] E. Nepovinnikh, T. Eerola, V. Biard, P. Mutka, M. Niemi, M. Kunasranta, H. Kälviäinen, Sealid: Saimaa ringed seal re-identification dataset, *Sensors* 22 (2022) 7602.
- [43] C. Trotter, G. Atkinson, M. Sharpe, K. Richardson, A. S. McGough, N. Wright, B. Burville, P. Berggren, NDD20: A large-scale few-shot dolphin dataset for coarse and fine-grained categorisation, arXiv preprint arXiv:2005.13359 (2020).
- [44] O. Wahltinez, S. J. Wahltinez, An open-source general purpose machine learning framework for individual animal re-identification using few-shot learning, *Methods in Ecology and Evolution* 15 (2024) 373–387.
- [45] L. Wang, R. Ding, Y. Zhai, Q. Zhang, W. Tang, N. Zheng, G. Hua, Giant panda identification, *IEEE Transactions on Image Processing* 30 (2021) 2837–2849.
- [46] WildMe, Beluga ID 2022, 2022. URL: <https://lila.science/datasets/beluga-id-2022>.
- [47] Botswana Predator Conservation Trust, *Panthera pardus* CSV custom export. Retrieved from African Carnivore Wildbook 2022-04-28, 2022.
- [48] P. Yu, T. Burghardt, A. W. Dowsey, N. W. Campbell, Multicamcows2024—a multi-view image

- dataset for ai-driven holstein-friesian cattle re-identification on a working farm, arXiv preprint arXiv:2410.12695 (2024).
- [49] M. Zuerl, R. Dirauf, F. Koeferl, N. Steinlein, J. Sueskind, D. Zanca, I. Brehm, L. v. Fersen, B. Eskofier, Polarbearvidid: A video-based re-identification benchmark dataset for polar bears, *Animals* 13 (2023) 801.
 - [50] S. Zuffi, A. Kanazawa, T. Berger-Wolf, M. J. Black, Three-D safari: Learning to estimate zebra pose, shape, and texture from images "In the wild", in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 5359–5368.
 - [51] L. Pícek, J. Straka, M. Jířík, E. Belotti, M. Duřá, J. Krausová, M. Bojda, V. Cermak, L. Buřka, R. Dvořák, et al., Czechlynx: A dataset for individual identification and pose estimation of the eurasian lynx, *Scientific Data* (2026).
 - [52] D. Biffi, M. R. Tucker, A. Ackel, D. A. Williams, Identification of individual texas horned lizards (*phrynosoma cornutum*) using genotypes and ventral spot patterns, *Ecology and Evolution* 15 (2025) e71167.
 - [53] A. Rueda-Toicen, A. A. Martin, D. Morozov, M. Mahmood, A. Schild, S. Dayani, D. Panza, G. de Melo, Are we recognizing the jaguar or its background? a diagnostic framework for jaguar re-identification, arXiv preprint arXiv:2604.09690 (2026).
 - [54] L. Pícek, T. Haucke, L. Adam, E. Nepovinnikh, L. Otashvili, K. Papafitsoros, T. Berger-Wolf, M. B. Brown, T. Burghardt, V. Cermak, et al., Centering ecological goals in automated identification of individual animals, arXiv preprint arXiv:2604.20626 (2026).
 - [55] M. Elbassat, S. Elkerdany, M. Elbially, G. Abouelhamd, J. Ghoneim, A. Elkady, M. Elboraay, N. Semenova, Calibrated similarity and graph clustering for open-set animal re-identification, in: *Working Notes of CLEF 2026 - Conference and Labs of the Evaluation Forum*, 2026.
 - [56] M. Yesharim, Y. Ram, Global embeddings, local feature matching, and adaptive validation overfitting in animal re-identification, in: *Working Notes of CLEF 2026 - Conference and Labs of the Evaluation Forum*, 2026.
 - [57] B.-A. Chen, P.-H. Shen, Y.-Y. Hsu, Species-specific multi-stage clustering for animal re-id, in: *Working Notes of CLEF 2026 - Conference and Labs of the Evaluation Forum*, 2026.
 - [58] E. S. Smith, A. Miyaguchi, S. Palamari, D. Evangelista, Ds@gt arc at animalclef 2026: Species-aware graph construction for multi-species animal re-identification, in: *Working Notes of CLEF 2026 - Conference and Labs of the Evaluation Forum*, 2026.
 - [59] G. Colling, Anchored graph clustering with local feature matching for animalclef 2026, in: *Working Notes of CLEF 2026 - Conference and Labs of the Evaluation Forum*, 2026.
 - [60] B. Liu, Z. Han, Z. Liang, Species-aware metric graph clustering for animal re-identification and discovery at animalclef 2026, in: *Working Notes of CLEF 2026 - Conference and Labs of the Evaluation Forum*, 2026.
 - [61] H. N. Rofiq, Lineage-aware markov partition refinement for multi-species animal re-identification, in: *Working Notes of CLEF 2026 - Conference and Labs of the Evaluation Forum*, 2026.
 - [62] X. Fan, Rapid exploration of foundation models and supervised contrastive learning for open-set individual animal discovery, in: *Working Notes of CLEF 2026 - Conference and Labs of the Evaluation Forum*, 2026.
 - [63] W. Chen, D. Cao, A prototype-constrained discovery clustering system for animalclef 2026, in: *Working Notes of CLEF 2026 - Conference and Labs of the Evaluation Forum*, 2026.
 - [64] F. Ye, X. Wang, K. Li, Z. Han, A species-wise clustering pipeline with salamander edge reliability learning for animalclef 2026, in: *Working Notes of CLEF 2026 - Conference and Labs of the Evaluation Forum*, 2026.
 - [65] D. Swaroop, V. Thenkanidiyoor, A. D. Dileep, Matching the baseline, and why that is worth writing about: A post-mortem of an animalclef 2026 re-identification pipeline, in: *Working Notes of CLEF 2026 - Conference and Labs of the Evaluation Forum*, 2026.
 - [66] S. Pourazari, Exploration of unsupervised segmentation as a foundation for vision graph neural networks in individual animal re-identification, in: *Working Notes of CLEF 2026 - Conference and Labs of the Evaluation Forum*, 2026.