

Overview of AnimalCLEF 2025: Recognizing Individual Animals in Images

Lukáš Adam^{1,*}, Kostas Papafitsoros², Roman Kovář³, Vojtěch Čermák⁴ and Lukas Pícek^{5,6}

¹Faculty of Electrical Engineering, University of West Bohemia, Czech Republic

²Queen Mary University of London, London, United Kingdom

³Faculty of Agrobiological Food and Natural Resources, Czech University of Life Science, Czech Republic

⁴Faculty of Electrical Engineering, Czech Technical University, Czech Republic

⁵Department of Cybernetics, Faculty of Applied Sciences, University of West Bohemia, Czech Republic

⁶Inria, LIRMM, University of Montpellier, France

Abstract

The first edition of the individual animal identification challenge, AnimalCLEF 2025, organized within LifeCLEF, advances the field of animal re-identification using computer vision and machine learning. Building on the WildlifeReID-10k dataset and incorporating new data, AnimalCLEF 2025 challenges participants to recognize individual animals from images for three species: lynxes, salamanders and sea turtles. The mix of species with different image capture conditions attempts to make the submitted prediction models generalizable to unseen species. The competition attracted 270 participants across 230 teams, with 136 outperforming the provided baseline based on MegaDescriptor. 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, open-set recognition, animals, re-identification, individual identification, machine learning, computer vision, classification

1. Introduction

Animal re-identification is crucial for studying various aspects of wildlife, including disease monitoring and control [1], assessing an animal's role in the ecosystem [2], tracking invasive species [3], and evaluating human impact on habitats and ecological restoration efforts [4]. Despite variations in how animal re-identification is defined and approached across studies, its objective remains consistent. The primary aim is to accurately and efficiently identify individual animals within a species by recognizing unique traits, such as markings, patterns, or other distinctive features. Automating this process allows for collecting detailed, large-scale data on population dynamics, migration routes, habitat use, and behaviors. This data empowers researchers to perform research on survival rate estimation [5], behavioural studies [6], injury and disease monitoring [7], measuring ecotouristic pressure [8] or inferring population distribution [9]. Such information is essential for identifying biodiversity threats and crafting evidence-based conservation strategies. Despite great progress in machine learning, many models for animal re-identification still struggle with overfitting. They often focus too much on background details - like lighting or landscape - rather than the animal's unique features. As a result, these models perform well in familiar environments but fail in new or varied conditions. This limitation makes it harder to accurately identify animals in diverse habitats, reducing the models' usefulness in real-world conservation efforts. Improving their ability to generalize is essential for more reliable animal identification.

CLEF 2025 Working Notes, 9 – 12 September 2025, Madrid, Spain

*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); 0000-0003-4411-9786 (R. Kovář);

0000-0002-9699-9220 (V. Čermák); 0000-0002-6041-9722 (L. Pícek)



© 2025 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 2025.

To address these issues and enable continual incremental improvements in animal re-identification, we organize the AnimalCLEF research competition. This first edition, which took part in the LifeCLEF 2025 [10] and the FGVC12 workshop at CVPR 2025, will hopefully start continual annual challenges, similar to other LifeCLEF challenges. This challenge builds on the vast animal re-identification WildlifeReID-10k [11] dataset with over 10k individual animals and provides new images of three species: lynxes, salamanders and sea turtles (see Figure 1). Given a split to a database and a query set, for each query image, the participants need to determine whether the individual is present in the database. If it is present, they need to determine its identity. Therefore, AnimalCLEF 2025 uses the open-set split instead of the classical closed-set split, which is unrealistic for wild animal populations. To provide a starting point for participants, we provided a baseline at the challenge start and then its improved version in the middle of the challenge. The initial baseline was the deep model MegaDescriptor [12] model, while the second baseline combined it with the local feature model ALIKED [13] by WildFusion [14]. We consider this competition a big success, as the state-of-the-art baseline, published in the same year as the challenge, was outperformed by 88 teams, bringing new ideas to the field of animal re-identification.

By promoting innovation in this field, AnimalCLEF 2025 aims to bridge further the gap between computer vision capabilities and practical conservationist needs, potentially impacting areas ranging from biodiversity research to habitat protection and beyond.

2. Challenge Description

2.1. Dataset

The AnimalCLEF 2025 dataset builds upon the previous editions of LifeCLEF [15–17] and the WildlifeReID-10k dataset [11]. The WildlifeReID-10k dataset was created by using the Python package wildlife-datasets [12] by combining the following datasets [18–47]. All of the images are cropped to contain the part used to recognise 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.

In addition to the training set, we assembled a test dataset with images never seen and shared on the internet, including three different species, i.e., *Lynx lynx*, *Salamandra salamandra*, and *Caretta caretta*. The lynx data is part of the newly released dataset, CzechLynx [48]. The turtle data is from an unreleased part of the SeaTurtleID2022 dataset [19]. The salamander data has not yet been released. All three datasets were split into the database and query sets, which are briefly described below, and their statistics in detail are listed in Table 1.

Table 1

AnimalCLEF 2025 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 database.

Subset	Number of Images	Number of Individuals		
		Total	Known	Unknown
WildlifeReID-10k (<i>Training</i>)	140,488	10,772	10,772	-
<i>Lynx lynx</i>	3,903	99	99	-
↳ <i>Database</i>	2,957	77	77	-
↳ <i>Query</i>	946	41	19	22
<i>Salamandra salamandra</i>	2,077	646	646	-
↳ <i>Database</i>	1,388	587	587	-
↳ <i>Query</i>	689	108	59	49
<i>Caretta caretta</i>	9,229	488	488	-
↳ <i>Database</i>	8,729	438	438	-
↳ <i>Query</i>	500	100	50	50

The subset with **Eurasian lynxes** (*Lynx lynx*) individuals includes real camera-trap photographs and synthetic samples, organized around three computer-vision tasks: individual identification, pose estimation, and instance segmentation. It spans more than 15 years and covers two regions in Central Europe (Southwest Bohemia and the Western Carpathians) and contains around 37k manually verified and labeled images. The images used for the AnimalCLEF 2025 challenge all contain real animals. The splits were provided by the authors [48].

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 without being marked.

The subset with **sea turtles** (*Caratta carreta*) individuals contains 8,729 photographs of 438 individuals collected within 13 years at 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 [19]. As the query set, we additionally published 500 images of 100 individuals, 50 of which were already present in the database and 50 were new.

2.2. Task and Evaluation Protocol

For all query images, the competitors needed to decide whether the depicted individual is known (present in the database). If yes, they additionally needed to specify the identity (label) of the individual. Therefore, the task was essentially open-set. The training set was completely optional, allowing trading fast training for accuracy.

The evaluation was performed by accuracy on known and unknown samples. The accuracy was further balanced with respect to datasets and individuals, preventing overrepresented individuals from

having too big an impact on the solution. Mathematically:

$$\begin{aligned} \text{BAKS} &= \frac{1}{D} \sum_{d=1}^D \frac{1}{|C_d|} \sum_{c \in C_d} \frac{1}{|I^{d,c}|} \sum_{i \in I^{d,c}} \mathbf{1}(\hat{y}_i = y_i), \\ \text{BAUS} &= \frac{1}{D} \sum_{d=1}^D \frac{1}{|C_d^*|} \sum_{c \in C_d^*} \frac{1}{|I^{d,c}|} \sum_{i \in I^{d,c}} \mathbf{1}(\hat{y}_i = \text{new}), \end{aligned} \quad (1)$$

where $\mathbf{1}$ is the indicator (0/1) function, $D = 3$ the number of datasets, C_d (or C_d^*) the *known* (or *unknown*) individuals (classes) in dataset d and $I^{d,c}$ indices of all images of individual c in the query set of the dataset d . Finally, y_i and $\hat{y}_i$ are the true and predicted identities of the i -th image.

To evaluate overall performance with a single score, we combine the closed-set (BAKS) and open-set (BAUS) metrics using their geometric mean:

$$\text{normalized accuracy} = \sqrt{\text{BAKS} \cdot \text{BAUS}}. \quad (2)$$

The geometric mean is preferred over the arithmetic mean, as it penalises extreme imbalances, e.g., a trivial classifier predicting only new individuals yields $\text{BAKS} = 0\%$ and $\text{BAUS} = 100\%$, resulting in a misleading 50% arithmetic mean but a more appropriate geometric mean of 0%.

2.3. Baselines

To enable an easier start for all participants and straightforward model evaluation, we provide two baselines based on the pre-trained models. The first one was the [MegaDescriptor-L-384](#) [12], which is a Swin [49] model trained on a dataset similar to the training set of the AnimalCLEF 2025 challenge. The second one is WildFusion [14], which combines MegaDescriptor with the local-feature extractor ALIKED [13] by a calibration process. Neither of these models was fine-tuned on the database.

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

2.4. Timeline

The AnimalCLEF 2025 competition was launched on February 28, 2025, and was promoted through the LifeCLEF [50], Kaggle, and FGVC challenge web pages, inviting participants to register. The competition ran for approximately three months, with the final submission deadline on May 19. 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 two submissions, which were evaluated on the private query set. The participants were scored and ranked based on the better of the two selected solutions.

2.5. Working Notes

Participants were strongly encouraged to submit both their code and a detailed technical report (Working Notes) to ensure their results can be fully reproduced. All the submitted Working Notes underwent thorough review and were given complex feedback by 2-3 experts with extensive publication records in Computer Vision and Machine Learning. 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

270 participants in 230 teams participated, with 136 teams outperforming the MegaDescriptor baseline (with normalized accuracy 0.309) and 88 teams outperforming the MegaDescriptor+ALIKED baseline (with normalized accuracy 0.443). The provided baseline was very good at detecting new individuals (BAUS of 0.992), while it had problems correctly classifying the known individuals (BAKS of 0.209).

From the Kaggle platform, we were able to extract all submissions but not the two final submissions that the participants selected for the final evaluation as described in Section 2.4. To analyze the participants' performance for this paper, we selected their submission with the highest public score and evaluated it on the private leaderboard. Therefore, the analysis presented here may slightly differ from the official leaderboard. The ranking of teams is taken from the official ranking.

Figure 2 (top) shows the results of the best 50 participants. All results are shown on the private dataset. The dash-dotted lines depict the normalized accuracy on individual datasets, while the full blue is the overall normalized accuracy computed as their mean. The normalized accuracy on sea turtles (dash-dotted orange) is relatively stable for all teams. The ranking of teams was most influenced by their performance on lynxes (dash-dotted red) and salamanders (dash-dotted black). Interestingly, the second-ranked team (*webmaking*) achieved the best performance on salamanders and sea turtles but was ranked 25 on lynxes. The winning team (*DataBoom*) achieved a more consistent performance on all datasets. If the results of these two best-performing teams were combined, the accuracy would increase from 0.713 (*DataBoom*) or 0.675 (*webmaking*) to 0.734.

Figure 2 (bottom) shows BAKS and BAUS averaged over all datasets. These numbers show that all teams achieved very high BAUS, that is, they predicted new individuals very well. We also show precision on unknown samples, defined as the ratio of correctly predicted unknowns over all predicted unknowns. Since the precision is between 0.5 and 0.6 and BAKS between 0.9 and 1.0, that means that all participants predicted approximately twice the number of unknown individuals than the true number of unknown individuals and that these predictions contained almost all unknown individuals.

4. Participants and Methods

We have received eight submissions, all of which outperformed the baseline MegaDescriptor, and six outperformed the WildFusion baseline. Six submissions were accepted into the proceedings. Details of the best methods and systems used are synthesised below and further developed in participants' working notes [51–56].

Team **DataBoom** [51] (Top 1) improved on WildFusion by using multiple local feature-matching techniques and by replacing MegaDescriptor with MiewID. They used MiewID as a preselected which did not enter the final score computation. The scores from local feature-matching techniques are combined by averaging. The authors also argued that some photos must be properly segmented to remove artifacts such as human fingers for salamanders.

Team **webmaking** [52] (Top 2) used two streams: (i) WildFusion, which provides calibrated global–local similarity scores, followed by two reranking methods applied to the top-20 images, and (ii) an XGBoost model trained on concatenated features from MegaDescriptor and MIEW-ID. The authors tuned the weights of both streams for all three datasets and provided additional experiments with preprocessing contour masks, image alignment and background blurring, which, however, did not bring a significant gain in accuracy.

Team **Tim Riggins** [53] (Top 7) combined the WildFusion baseline with the gradient boosting method LightGBM [57], which is an inexact fast approximation of the conventional Gradient Boosting Decision Tree (GBDT) method [58]. The authors did not perform any fine-tuning of the feature extractor, making their algorithm efficient.

Team **already mygo** [54] (Top 34) used a similar approach to the WildFusion baseline, thus a combination of deep method (Swin) with local features (DISK and ALIKED) combined via the WildFusion method. The authors used fine-tuning of the feature extractor on the database and performed an analysis of the thresholding parameter used to classify new individuals.

Team **dongyeonkim10** [55] (Top 44) used data preprocessing techniques and then averaged the WildFusion score with EVA02 [59] score. The authors also investigated inference-time augmentations and included techniques which did not work for them.

Team **DS@GT LifeCLEF** [56] (Top 103) compares features extracted by DINOv2 and MegaDescriptor-L, reprojected using triplet loss, and classified using a k-nearest neighbors (kNN) classifier. Since they fine-tuned the baseline MegaDescriptor, it outperforms it, but falls short of top leaderboard entries.




Figure 2: Private Leaderboard of the AnimalCLEF 2025 competition. We report the performance of the top 50 (out of 230) teams. The track metrics are the normalized accuracy split on individual datasets (top subfigure) and BAKS, BAUS and precision combined on all datasets (bottom subfigure).

5. Conclusions and Discussion

This paper presents an overview and evaluation of the results of the AnimalCLEF 2025 challenge, organized in conjunction with the LifeCLEF and CVPR-FGVC12 workshops. The challenge introduced a multi-species, open-set re-identification task using the WildlifeReID-10k dataset for training and new identities from three species, e.g., *Caretta caretta*, *Salamandra salamandra*, and *Euroasian lynx*, for testing. Participants were tasked with developing a model that automatically determines whether a query image belongs to a known individual in the gallery set or represents a new, unseen individual. The task emphasized robustness under real-world conditions, including background clutter, low sample counts, and interspecific variation. The competition attracted over 230 teams and demonstrated a wide variety of methodological approaches. Notably, global feature extractors, such as MegaDescriptor, remained central, while a resurgence in local matching and calibrated score fusion strategies highlighted the importance of fine-grained cues. Here are the key takeaways from this edition:

- **Global-local fusion is the new baseline.** Nearly all top-performing teams employed hybrid methods combining global embeddings (e.g., MegaDescriptor, MIEW) with local feature matchers (e.g., ALIKED, DISK, SuperPoint). These were fused using either WildFusion or custom weighted schemes, leading to significant gains over global-only systems.
- **Score calibration is essential in open-set recognition.** Teams that succeeded in balancing BAKS and BAUS scores commonly applied threshold-based novelty detection combined with isotonic or learned calibration of similarity scores. Without this, methods either overpredicted new individuals or failed to detect them altogether.
- **Diversity in global embedding sources boosts generalization.** Teams leveraging multiple descriptor models (e.g., Swin-L from MegaDescriptor and CNN-based MIEW or EVA02) demonstrated that architectural diversity leads to more robust predictions across species and conditions.
- **Meta-learning and ensemble stacking improve robustness.** The runner-up solution implemented a meta-algorithm with staged fusion: global-local WildFusion scores fed into an XGBoost classifier, followed by ArcFace fine-tuning per species. This approach showed cumulative improvements at each stage.

Directions for Future Editions. The first edition of AnimalCLEF established a strong foundation for evaluating open-set animal re-identification. Looking ahead, we identify several key directions that could enrich future editions and better reflect real-world biodiversity monitoring needs:

1. **Temporal robustness and longitudinal tracking.** Extending the task to multi-year observations (e.g., from juvenile to adult) would test models' ability to handle appearance changes over time; a critical requirement in wildlife research.
2. **Species-agnostic or cross-species modeling.** Unified models that generalize across multiple or unseen species would reflect real-world deployments, especially in diverse environments.
3. **Identity and category discovery.** Introducing unsupervised or semi-supervised settings where new individuals or even species must be grouped or discovered would push the boundaries of open-world recognition.
4. **Multimodal and metadata-aware tasks.** Although image-only systems dominated in 2025, future tracks could explicitly encourage the use of spatial, temporal, and ecological metadata to enrich recognition and enable more contextual reasoning.

We believe these extensions would not only stimulate innovation in machine learning but also bring AnimalCLEF closer to practical conservation and ecological monitoring scenarios.

6. 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

LA has been supported by the Ministry of Education, Youth and Sports of the Czech Republic under the project SGS-2024-017.

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] 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. URL: <https://doi.org/10.1002/ecy.3027>.
- [6] 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. URL: <https://doi.org/10.1016/j.anbehav.2022.05.006>.
- [7] K. Papafitsoros, A. Panagopoulou, G. Schofield, Social media reveals consistently disproportionate tourism pressure on a threatened marine vertebrate, *Animal Conservation* 24 (2021) 568–579. URL: <https://doi.org/10.1111/acv.12656>.
- [8] K. Papafitsoros, L. Adam, G. Schofield, A social media-based framework for quantifying temporal changes to wildlife viewing intensity, *Ecological Modelling* 476 (2023) 110223. URL: <https://doi.org/10.1016/j.ecolmodel.2022.110223>.
- [9] 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. URL: <https://doi.org/10.3389/fmars.2021.671061>.
- [10] 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.
- [11] 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.
- [12] 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.
- [13] X. Zhao, X. Wu, W. Chen, P. C. Y. Chen, Q. Xu, Z. Li, Aliked: A lighter keypoint and descriptor extraction network via deformable transformation, *IEEE Transactions on Instrumentation & Measurement* 72 (2023) 1–16. URL: <https://arxiv.org/pdf/2304.03608.pdf>. doi:10.1109/TIM.2023.3271000.
- [14] 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.
- [15] T. Lorieul, E. Cole, B. Deneu, M. Servajean, A. Durso¹¹, I. Bolon, H. Glotin, R. Planqué, R. R. de Castaneda, W.-P. Vellinga, et al., Overview of lifeclef 2021: An evaluation of machine-learning based species identification and species distribution prediction, in: *Experimental IR Meets Multilinguality, Multimodality, and Interaction: 12th International Conference of the CLEF Association, CLEF 2021, Virtual Event, September 21–24, 2021, Proceedings*, volume 12880, Springer Nature, 2021, p. 371.
 - [16] A. Joly, H. Goëau, S. Kahl, B. Deneu, M. Servajean, E. Cole, L. Pícek, R. R. De Castaneda, I. Bolon, A. Durso, et al., Overview of lifeclef 2020: a system-oriented evaluation of automated species identification and species distribution prediction, in: *International Conference of the Cross-Language Evaluation Forum for European Languages*, Springer, 2020, pp. 342–363.
 - [17] 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.
 - [18] L. Adam, K. Papafitsoros, C. Jean, A. Rees, Exploiting facial side similarities to improve ai-driven sea turtle photo-identification systems, *bioRxiv* (2024) 2024–09.
 - [19] 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.
 - [20] 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. URL: <https://www.sciencedirect.com/science/article/pii/S0168169921001514>.
 - [21] 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. URL: https://openaccess.thecvf.com/content_ICCV_2017_workshops/w41/html/Andrew_Visual_Localisation_and_ICCV_2017_paper.html.
 - [22] 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. URL: <https://ieeexplore.ieee.org/abstract/document/7532404>.
 - [23] 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. URL: <https://ieeexplore.ieee.org/document/9096922>.
 - [24] Local Ocean Conservation, Turtle recall: Conservation challenge, 2021. URL: <https://zindi.africa/competitions/turtle-recall-conservation-challenge>.
 - [25] 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 (ISCMCI)*, IEEE, 2020, pp. 224–228.
 - [26] 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.
 - [27] 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.
 - [28] L. Fu, G. He, Cow dataset, 2021. URL: <https://doi.org/10.6084/m9.figshare.16879780.v1>.
 - [29] 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). URL: <https://arxiv.org/abs/2105.01938>.
 - [30] 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.
- [31] 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. URL: <https://www.int-res.com/abstracts/esr/v7/n1/p39-53/>.
 - [32] 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.
 - [33] 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.
 - [34] 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.
 - [35] T.-Y. Lin, Cat individual images, 2020. URL: <https://www.kaggle.com/datasets/timost1234/cat-individuals>.
 - [36] V. Miele, G. Dussert, B. Spataro, S. Chamaillé-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.
 - [37] 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.
 - [38] E. Nepovinnikh, T. Eerola, V. Biard, P. Mutka, M. Niemi, M. Kunnasranta, H. Kälviäinen, Sealid: Saimaa ringed seal re-identification dataset, *Sensors* 22 (2022) 7602.
 - [39] 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. URL: <https://www.aaai.org/ocs/index.php/SSS/SSS17/paper/viewPaper/15245>.
 - [40] 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). URL: <https://arxiv.org/abs/2005.13359>.
 - [41] 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.
 - [42] 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. URL: <https://ieeexplore.ieee.org/document/9347819>.
 - [43] WildMe, Beluga ID 2022, 2022. URL: <https://lila.science/datasets/beluga-id-2022>.
 - [44] WildMe, Leopard ID 2022, 2022. URL: <https://africancarnivore.wildbook.org/>.
 - [45] 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).
 - [46] 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.
 - [47] 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. URL: <https://ieeexplore.ieee.org/document/9010937>.
 - [48] L. Pícek, E. Belotti, M. Bojda, L. Buřka, V. Cermak, M. Dula, R. Dvorak, L. Hrdy, M. Jirik, V. Kocourek, J. Krausova, J. Labuda, J. Straka, L. Toman, V. Trulík, M. Vana, M. Kutal, Czechlynx: A dataset for individual identification and pose estimation of the eurasian lynx, *arXiv preprint arXiv:1804.07177* (2025).
 - [49] Z. Liu, H. Hu, Y. Lin, Z. Yao, Z. Xie, Y. Wei, J. Ning, Y. Cao, Z. Zhang, L. Dong, et al., Swin

- transformer v2: Scaling up capacity and resolution, in: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2022, pp. 12009–12019.
- [50] A. Joly, L. Picek, S. Kahl, H. Goëau, L. Adam, C. Botella, M. Servajean, D. Marcos, C. Leblanc, T. Larcher, et al., Lifeclef 2025 teaser: Challenges on species presence prediction and identification, and individual animal identification, in: European Conference on Information Retrieval, Springer, 2025, pp. 373–381.
 - [51] S. Lansikh, G. Demidov, D. Dinmuhametov, A. Khlopotnukh, K. Bogdan, R. Pakhomov, Matchers with weights, in: Working Notes of CLEF 2025 - Conference and Labs of the Evaluation Forum, 2025.
 - [52] N. Semenova, Dual-stream meta-algorithm for animal re-id: Wildfusion global-local matching and xgboost on megadescriptor-miew embeddings, in: Working Notes of CLEF 2025 - Conference and Labs of the Evaluation Forum, 2025.
 - [53] S. Fedorchenko, S. Arefiev, Gradient boosting similarity in entity matching, in: Working Notes of CLEF 2025 - Conference and Labs of the Evaluation Forum, 2025.
 - [54] J. Tan, A. Wang, Open-set animal re-identification via multilevel feature fusion, in: Working Notes of CLEF 2025 - Conference and Labs of the Evaluation Forum, 2025.
 - [55] D. Kim, B. Park, H. Bae¹, S. Lee, C. Lee, Fusion of global and local descriptors with feature calibration for multi-species animal re-identification: Animalclef 2025, in: Working Notes of CLEF 2025 - Conference and Labs of the Evaluation Forum, 2025.
 - [56] A. Miyaguchi, C. Maruthaiyannan, C. R. Clark, Ds@gt animalclef: Triplet learning over vit manifolds with nearest neighbor classification for animal re-identification, in: Working Notes of CLEF 2025 - Conference and Labs of the Evaluation Forum, 2025.
 - [57] G. Ke, Q. Meng, T. Finley, T. Wang, W. Chen, W. Ma, Q. Ye, T.-Y. Liu, Lightgbm: A highly efficient gradient boosting decision tree, *Advances in neural information processing systems* 30 (2017).
 - [58] J. H. Friedman, Greedy function approximation: a gradient boosting machine, *Annals of statistics* (2001) 1189–1232.
 - [59] Y. Fang, Q. Sun, X. Wang, T. Huang, X. Wang, Y. Cao, Eva-02: A visual representation for neon genesis, *Image and Vision Computing* 149 (2024) 105171.