Chicks4FreeID: A Benchmark Dataset for Chicken Re-Identification

Daria Kern[1,2] * **Tobias Schiele**[1,2] * **Ulrich Klauck**[1,3] **Winfred Ingabire**[2]

[1]Aalen University, Germany
{daria.kern, tobias.schiele, ulrich.klauck}@hs-aalen.de
[2]Glasgow Caledonian University, United Kingdom
winfred.ingabire@gcu.ac.uk
[3]University of the Western Cape, South Africa

Figure 1: Excerpt from the Chicks4FreeID dataset.

Abstract

To address the need for well-annotated datasets in the field of animal re-identification, and particularly to close the existing gap for chickens, we introduce the Chicks4FreeID dataset. This dataset is the first publicly available re-identification resource dedicated to the most farmed animal in the world. It includes top-down view images of individually segmented and annotated chickens, along with preprocessed cut-out crops of the instances. The dataset comprises 1215 annotations of 50 unique chicken individuals, as well as a total of 55 annotations of 2 roosters and 2 ducks. In addition to re-identification, the dataset supports semantic and instance segmentation tasks by providing corresponding masks. Curation and annotation were performed manually, ensuring high-quality, nearly pixel-perfect masks and accurate ground truth assignment of the individuals using expert knowledge. Additionally, we provide context by offering a comprehensive overview of existing datasets for animal re-identification. To facilitate comparability, we establish a baseline for the re-identification task testing different approaches. Performance is evaluated based on mAP, Top-1, and Top-5 accuracy metrics. Both the data and code are publicly shared under a CC BY 4.0 license, promoting accessibility and further research. The dataset can be accessed at https://huggingface.co/datasets/dariakern/Chicks4FreeID and the code at https://github.com/DariaKern/Chicks4FreeID.

*contributed equally. Contact: Chicks4FreeID@dariakern.com

Submitted to the 38th Conference on Neural Information Processing Systems (NeurIPS 2024) Track on Datasets and Benchmarks. Do not distribute.

1 Introduction

1.1 Motivation

Chickens struggle to recognize other individuals after visible changes are applied to the comb or plumage [1]. Much like chickens are able to use visual cues to differentiate each other, artificial intelligence (AI) is capable of utilizing image or video inputs for re-identification purposes. AI-driven re-identification and tracking systems hold great potential for enhancing animal husbandry and livestock farming. These systems may allow for the observation of social structures and behavior, enhance welfare, and potentially lead to more efficient animal management with minimal disruption to the livestock [2]. They also may help assess health and well-being, i.e., by providing crucial traceability during disease outbreaks. Furthermore, they offer a cost-effective and non-invasive alternative to manual tagging methods.

Despite the significant potential, there is a notable gap in publicly available datasets for such technologies, especially for chickens — the most farmed animal globally. Remarkably, to our knowledge, no publicly available dataset for chicken re-identification exists, highlighting an urgent need for development in this field. Public datasets for the task of individual animal re-identification in general are scarce [3, 2]. In particular, well-annotated datasets [4]. The practice of openly sharing data and code should be encouraged to enhance result comparability, yet not all research data are currently made public. In their work, [5] emphasize the importance of creating and sharing publicly available and well-annotated benchmark datasets for the task of animal re-identification.

Establishing a benchmark dataset involves evaluating how well existing methods solve the dataset. The reported metrics serve as a baseline for future researchers to report their improvements. Given the diverse nature of research, it is important for the baseline to cover common approaches and common metrics. This ensures that the achievements of future researchers can be effectively compared, facilitating a standardized assessment of advancements in the field.

1.2 Contribution

We address the existing gap and present our Chicks4FreeID dataset, which does not only support the task of re-identification but also semantic and instance segmentation. We make this thoroughly documented dataset freely accessible to the research community and the public. The dataset includes 54 individuals, of which 50 are chickens. Each occurrence is nearly pixel-perfectly segmented, resulting in 1270 instance masks. Based on the cut-out crops of 1215 chicken instance masks, we provide an initial baseline for the task of closed set re-identification. This allows the research community to compare their methods and results effectively. In summary:

i We provide a comprehensive overview of publicly available datasets for animal re-identification.

ii We introduce the first publicly available dataset for chicken re-identification.

iii We establish a baseline for closed set re-identification on the introduced dataset.

2 Related work

2.1 Animal re-identification

Animal re-identification, the task of identifying individual animals within one (or sometimes several) species, finds applications in various fields. Particularly in wildlife conservation efforts, where monitoring endangered species is crucial [6–9]. But also in livestock management, notably cattle [10–14] and yak [15]. Honeybees [4] and bumblebees [16, 17] have also been subject to investigation.

Re-identification falls into one of two categories: closed set and open set. In closed set re-identification, all individuals are known from the beginning, and those to be identified can be matched with identities of a predefined set. In open set re-identification, the identity of the individual in question may not necessarily be part of a predefined set. It is possible to encounter completely new, undocumented individuals. Such individuals must be annotated as a new identity and, upon subsequent encounters, accurately matched.

While facial recognition is a prevalent method for re-identifying humans [18], the faces of animals can likewise serve as a means to re-identify individuals, as has previously been demonstrated for rhesus macaque [19], chimpanzee [20], cats [21], lions [22], dogs [23], giant pandas [8] and red pandas [9]. However, animals frequently exhibit more distinctive visual traits beyond their faces. For example, natural markings such as stripes [24–27] and scale patterns [28] can serve as prominent identifiers. But also specific body parts can contribute to distinguishing individuals, such as the fins of dolphins [29] and sharks [30]. Similarly to how fingerprints differentiate humans, nose prints of dogs have been utilized to uniquely identify individual dogs [31]. Conversely, little inter-individual variability poses a challenge to the re-identification task. Species exhibiting minimal or subtle visual distinctions between individuals are, for instance, (polar) bears [32, 33] or elephants [34]. Visual traits play a pivotal role in animal re-identification within computer vision, serving as essential markers for distinguishing individuals. However, the task is complex and extends beyond mere visual cues. Factors such as lighting, perspective, body changes over time, and partially obscured body parts pose additional challenges [5].

To further advance the field and aid the research community, [35] released the WildlifeDatasets toolkit - an open-source toolkit for animal re-identification. It gathers publicly available animal re-identification datasets in one place, in an effort to make them more easily accessible and to improve usability. Included are various tools, i.e., for data handling and processing, algorithms relevant to the task of re-identification, pretrained models, as well as evaluation methods. Therewith, they address the prevailing absence of standardization across the literature and facilitate comparability and reproducibility of results. Within their work, they also introduce a new state-of-the-art, the MegaDescriptor, notably the first foundation model for animal re-identification. Likewise, [36] present an open-source re-identification method initially developed for sea stars, which was successfully extended to seven mammalian species without adjustments. They also report state-of-the-art results. Moreover, [37] introduced Tri-AI, a system designed for the rapid detection, identification, and tracking of individuals from a wide range of primate species. The system is capable of processing both video footage and still images. The task of re-identification is closely related to tracking, where individuals are detected and tracked across various video frames. During tracking, individuals often need to be re-identified after leaving and re-entering the field of vision.

2.2 Re-identification datasets

A review of existing resources revealed fewer than 40 publicly available datasets for animal re-identification. This leads to the conclusion that a significant number of animal species are not yet covered, including chickens. Birds in general seem to be underrepresented in this domain, with only a couple of datasets available [38, 39]. In fact, a noticeable focus lies on marine life [40–50]. However, cattle are the most frequently featured species [11, 51–55], with much of the data collected by the same group of researchers.

Table 1 provides a summary of the publicly accessible datasets found, arranged by year. Each entry details the name of the dataset ("Dataset"), the associated publication ("Publ."), and species focus ("Species"). "IDs" denotes the number of unique identities present within the dataset. Additionally, the total number of annotated animal instances within all images of each dataset is noted ("Annot."). An indication(*) of whether the data was derived from video sources is given as well. For ease of access, a direct link to each dataset is provided ("Avail. at"). Although all of the datasets are publicly accessible, some are released under licenses that are relatively restrictive.

Table 1: Publicly available animal re-identification datasets, arranged by date of publication. An asterisk (*) marks data derived from video footage.

Year	Publ.	Dataset	IDs	Species	Annot.	Avail. at
	ours	Chicks4FreeID	50, 2, 2	chicken, duck, rooster	1215, 40, 15	[56]
2024	[28]	SeaTurtleID2022	438	sea turtle	8729	[40]
2023	[3]	Mammal Club (IISD)	218	11 terrestrial mammal species*	33612	[57]
2023	[58]	Multi-pose dog dataset	192	dog	1657	[59]
2023	[32]	PolarBearVidID	13	polar bear*	138363	[60]
2023	[36]	Sea Star Re-ID	39, 56	common starfish, Australian cushion star	1204, 983	[41]
2022	[61]	Animal-Identification-from-Video	58, 26, 9	pigeon*, pig*, Koi fish*	12671, 6184, 1635	[39]
2022	n.a.	Beluga ID	788	beluga whale	5902	[42]
2022	n.a.	Happywhale	15587	30 different species of whales and dolphins	51033	[43]
2022	n.a.	Hyiena ID	256	spotted hyena	3129	[62]
2022	n.a.	Leopard ID	430	African leopard	6805	[63]
2022	[64]	SealID	57	Saimaa ringed seal	2080	[44]
2022	[65]	SeaTurtleIDHeads	400	sea turtle	7774	[45]
2022	n.a.	Turtle Recall	100	sea turtle	2145	[46]
2021	[66]	Cow Dataset	13	cow	3772	[11]
2021	[13]	Cows2021	182	Holstein-Friesian cattle*	13784	[51]
2021	[67]	Giraffe Dataset	62	giraffe	624	[68]
2021	[8]	iPanda-50	50	giant panda	6874	[69]
2020	[26]	AAU Zebrafish Dataset	6	zebrafish*	6672	[70]
2020	[37]	Animal Face Dataset	1040	41 primate species	102399	[71]
2020	[24]	ATRW	92	Amur tiger*	3649	[72]
2020	[73]	Lion Face Dataset	94	lion	740	[22]
2020	[74]	NDD20	44, 82	bottlenose and white-beaked dolphin, white-beaked dolphin (underwater)*	2201, 2201	[47]
2020	[73]	Nyala Data	237	nyala	1942	[75]
2020	[14]	OpenCows2020	46	Holstein-Friesian cattle*	4736	[52]
2019	[76]	Bird individualID	30, 10 ,10	sociable weaver, great tit, zebra finch	51934	[38]
2019	[23]	Dog Face Dataset	1393	dog	8363	[77]
2018	[21]	Cat Individual Images	518	cat	13536	[78]
2018	[79]	Fruit Fly Dataset	60	fruit fly*	2592000	[80]
2018	n.a.	HumpbackWhaleID	5004	humpback whale	15697	[48]
2018	[19]	MacaqueFaces	34	rhesus macaque*	6280	[81]
2017	[12]	AerialCattle2017	23	Holstein-Friesian cattle*	46340	[53]
2017	[12]	FriesianCattle2017	89	Holstein-Friesian cattle*	940	[54]
2017	[25]	GZGC	2056	plains zebra and Masai giraffe	6925	[82]
2016	[20]	C-Tai	78	chimpanzee	5078	[83]
2016	[20]	C-Zoo	24	chimpanzee	2109	[83]
2016	[10]	FriesianCattle2015	40	Holstein-Friesian cattle*	377	[55]
2015	n.a.	Right Whale Recognition	447	North Atlantic right whale	4544	[49]
2011	[27]	StripeSpotter	45	plains and Grevy's zebra	820	[27]
2009	[84]	Whale Shark ID	543	whale shark	7693	[50]

3 The Chicks4FreeID dataset

3.1 Data

The Chicks4FreeID dataset contains top-down view images of individually segmented and annotated chickens, with some images also featuring roosters and ducks. Each image is accompanied by a color-coded semantic segmentation mask that classifies pixel values by animal category (chicken, rooster, duck) and background, as well as binary segmentation mask(s) for the animal instance(s) depicted. Additionally, the dataset includes preprocessed cut-out crops (detailed in Section 3.5) of the respective animal instances. Figure 2 gives a first overview of the dataset.

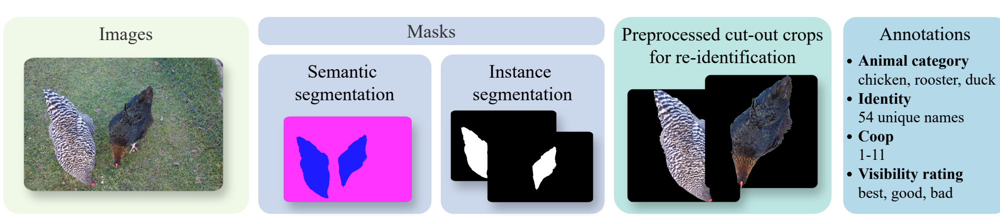

Figure 2: Dataset overview.

3.2 Collection

Various coops of private households were visited to photograph chickens. Among these coops, two additionally accommodate a rooster each, while another houses two ducks. A total of 677 images were captured using two similar models of cameras: the "Sony CyberShot DSC-RX100 VI" and the "Sony CyberShot DSC-RX100 I". The resolution of the images stands at 3648x5472 pixels. Every image includes at least one chicken, ensuring no images without chickens are part of the dataset. It was collected over the span of one year, however all images of a coop were shot within one day. In other words, all photos of a given individual were taken on the same day. The images were captured from a top-down view perspective, aiming to capture the plumage. The dataset is not limited to a single breed of chicken, ensuring a certain level of variability.

3.3 Annotation

We utilized Labelbox [85] under a free educational license for manual data annotation.

Instances and background For each animal instance appearing within an image, a segmentation was meticulously hand-crafted by a human annotator. No AI has been used during the annotation process to ensure high-quality, nearly pixel-perfect instance masks. The instance masks include the comb, head, beak, and plumage. Feet were excluded as rings could give away the identity. Feet and scattered feathers are considered part of the background, along with any visible objects or living beings that are not chickens, roosters, or ducks. Compared to conventional bounding boxes, instance masks offer the advantage of better supporting the subsequent re-identification process. The background can be easily removed as it might contain unwanted clues about the identity of the chickens. Furthermore, the provided masks render the dataset well-suited for instance segmentation tasks as well.

Animal categories Each instance of an animal was assigned to one of three animal categories. These are "chicken", "rooster", and "duck". Roosters and especially ducks serve as exceptions within the predominantly chicken-based collection. This characteristic potentially positions the dataset as a resource for anomaly detection as well.

Identities and coops The identities of the subjects were meticulously studied prior to photography, closely monitored throughout the image capture process, and ultimately assigned by a human annotator. The ground truth annotation was performed without the use of any algorithm. In cases where the human annotator could not assign an identity, the instance was labeled as identity "Unknown". It is essential to clarify that the label "Unknown" does not imply the presence of a new individual. Instead, it represents an unidentified individual from the closed set, more precisely, from the annotated coop. Each image contains one or more chickens, all of which are individually identified by their unique names. Roosters and ducks are each also uniquely named. Furthermore, each instance is explicitly annotated to indicate the specific coop to which it belongs.

Visibility rating Acknowledging varying visibility of the subjects (chickens, roosters, ducks) within the images, each appearance has been manually assigned a visibility rating, categorized as either "bad", "good", or "best". The "best" rating includes segmentation instances that fully display the subject from the desired top-down perspective, and those where only an insignificant part is missing, such as the very tip of the tail feathers. Instances that include only small parts of the subject and on which the subject is difficult to recognize fall under the "bad" rating. All remaining segmentation instances, that do not qualify as "bad" or "best", are rated as "good".

3.4 Composition

The dataset comprises a collection of 677 images, featuring a total of 50 distinct chicken, 2 rooster, and 2 duck identities distributed across 11 different coops. A total of 1270 instances were obtained by segmenting 1215 appearances (instances) of chickens, alongside 15 roosters and 40 ducks.

Each instance is of a certain animal category (“chicken”, “rooster”, “duck”) and was assigned the corresponding coop (1-11), visibility (“best”, “good”, “bad”) and identity (1 of 54 names or “Unknown”). It is important to mention that no “Unknown” instances are present in the “best” or “good” subset. The ground truth identity for all instances in these subsets is, therefore, known. Figure 3 illustrates the number of instances for each individual, as well as the visibility rating of the instances. It starts with the individual with the most instances in the “best” subset and is arranged in descending order. The most represented chicken in the “best” subset is Mirmir with 27 instances, whereas Isolde is the least represented chicken with 4 instances.

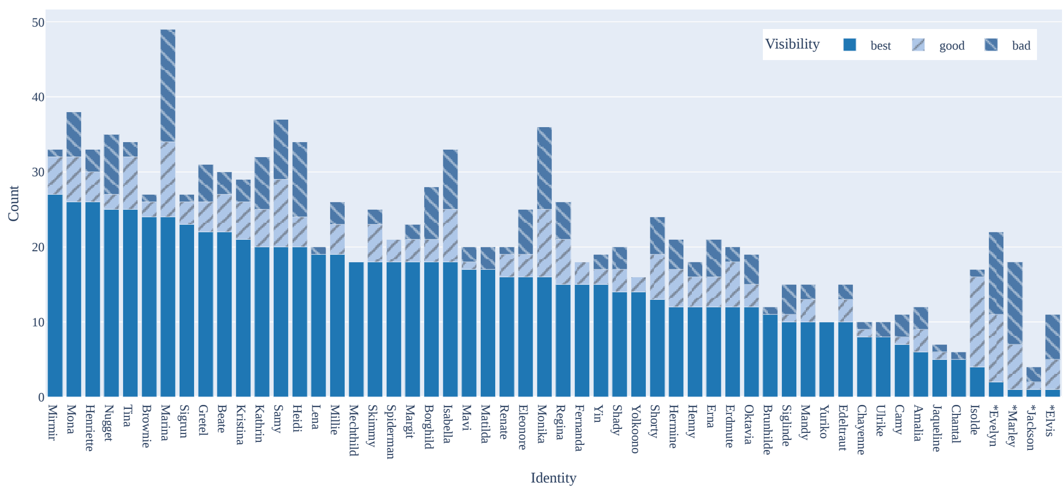

Figure 3: Visibility distributions for all instances of each individual. Ducks and roosters are marked with an asterisk (*).

3.5 Preprocessing

The following steps describe the preprocessing procedure to obtain the cut-out crops for the re-identification task. For all individuals captured in an image, a bounding box is created based on the instance masks. In the first step, both the image and the mask are cropped (to the area of interest contained in the bounding box) to focus solely on the individual (see Figure 4: Step 1). The cropped mask is then used to remove the background from the cropped image (Step 2). Finally, the resulting image is adjusted to a square shape for ease of use and consistency (Step 3). The resulting resolutions remain as is, with no resizing taking place.

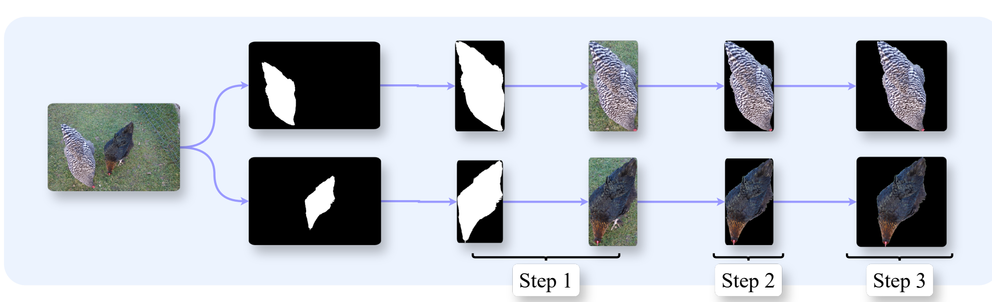

Figure 4: Data preprocessing pipeline for subsequent re-identification.

4 Experiments

4.1 Dataset, split and augmentation

For the closed set re-identification experiments, we utilize preprocessed cut-out crops as described in Section 3.5. To focus solely on all 50 chicken identities, the four identities of ducks and roosters were excluded. By removing instances of visibility level "good" and "bad", we ensure that only instances with "best" visibility are included. The utilized "best" subset does not contain any "Unknown" instances. The number of chicken instances contained in the "best" subset is 793.

The employed data is split into 630 train pairs and 163 test pairs of cut-out crops and the assigned identities. To ensure that the testing set does not introduce any new identities, we include all possible identities in the training set. For a fair evaluation on all identities, the train/test split is stratified, i.e., each identity has the same fixed percentage of its cut-out crops allocated to the test set. Consequently, identities with a higher total number of crops will contribute more to the test set compared to identities with fewer crops, ensuring proportional representation across all identities. The corresponding subset on Hugging Face is "chicken-re-id-best-visibility".

To avoid data leakage, it is important to apply data augmentation only after a train-test split is established. This ensures that augmented versions of the same original image do not appear in both sets. We dynamically apply the following data augmentation during training on the "chicken-re-id-best-visibility" subset: rotation, flip, RandAugment [86], and random color-jitter. No data augmentation is applied to the test set.

4.2 Baseline approaches

To establish a baseline for the closed set re-identification task, we test three different approaches on our dataset. Each approach involves two steps. First, a feature extractor generates embeddings for the cut-out crops. Second, the resulting feature vectors (embeddings) are then passed to a classifier to ultimately assign the identities. We test each approach with a variation of two classifiers: k-Nearest Neighbor (k-NN) and a linear classifier adapted from the Lightly library [87] (MIT License). All feature extractors were fed with images at an input resolution of 384 x 384 pixels and each approach was run three times. The baseline results were obtained on 64GB shared memory Apple M3 Max Chips (2023) running PyTorch 2.3.0 with MPS acceleration.

MegaDescriptor The employed MegaDescriptor-L-384 [35] (CC BY-NC 4.0 license [88]) is a state-of-the-art feature extractor for animal re-identification from the WildlifeDatasets toolkit (MIT license). It is based on the Swin Transformer architecture [89] and was pretrained on diverse datasets featuring various animal species. However, it has not been trained on chicken data and we did not fine-tune it either. A notable hyperparameter choice made by the MegaDescriptor-L384 authors is the ArcFace [90] loss function, which aims to aid in building meaningful embeddings. We selected the frozen MegaDescriptor-L-384 model over DINOv2 [91] and CLIP [92] due to its better performance on unseen animal domains, as reported by the authors. Their evaluation included cattle as an example of an unseen domain [35].

Swin Transformer We utilize the swin_large_patch4_window12_384 architecture [89] as implemented in [93]. We train it from scratch on the Chicks4FreeID dataset in a fully supervised manner. The training process and hyperparameters mirror those used to build the MegaDescriptor-L384, which also employs the same Swin Transformer architecture. Unlike the frozen MegaDescriptor-L384, which was trained on a variety of animal datasets, we now train the Swin architecture exclusively on our own dataset. The Swin Transformer itself is based on the Vision Transformer architecture.

Vision Transformer Finally, we employ the ViT-B/16 [94] architecture, as implemented in [95], and train it on the Chicks4FreeID dataset in a fully supervised manner with a simple cross-entropy loss. We adopted the effective hyperparameter settings as used in Lightly's benchmarks [87], including optimizer and scheduler choices, for our experiments. The difference between the Swin Transformer and the Vision Transformer lies in how they handle image data; the Swin Transformer uses a hierarchical structure with shifted windows to capture local and global features, while the Vision Transformer treats images as sequences of patches, relying on self-attention mechanisms throughout.

4.3 Evaluation

For all baselines, we provide three of the most common metrics for closed set animal re-identification. These are: mAP (mean Average Precision), Top-1 accuracy (ratio of correct predictions versus total predictions), and Top-5 accuracy (accuracy of the correct class being within the top 5 predictions) as implemented in TorchMetrics [96].

4.4 Baseline results and discussion

The results for all baseline approaches and the respective variations are summarized in Table 2. Overall, the experiments yield good results but still leave room for improvement.

Table 2: Baseline results for the closed set re-identification experiments. The highest scores for each metric are in blue.

Feature extractor	Training	Epochs	Classifier	mAP	Top-1	Top-5
MegaDescriptor [35]	pretrained, frozen	-	k-NN	0.649 ± 0.044	0.709 ± 0.026	0.924 ± 0.027
MegaDescriptor [35]	pretrained, frozen	-	linear	0.935 ± 0.005	0.883 ± 0.009	0.985 ± 0.003
Swin Transformer [89]	from scratch	200	k-NN	0.837 ± 0.062	0.881 ± 0.041	0.983 ± 0.010
Swin Transformer [89]	from scratch	200	linear	0.963 ± 0.022	0.922 ± 0.042	0.987 ± 0.012
Vision Transformer [94]	from scratch	200	k-NN	0.893 ± 0.010	0.923 ± 0.005	0.985 ± 0.019
Vision Transformer [94]	from scratch	200	linear	0.976 ± 0.007	0.928 ± 0.002	0.990 ± 0.012

Both the Swin Transformer and Vision Transformer architectures, when trained from scratch, outperformed the frozen MegaDescriptor model. Additionally, linear classifiers consistently outperformed k-NN classifiers. This indicates that performance scales with the level of supervision, which aligns with expectations.

The gap between the MegaDescriptor, a model from a different domain (trained on different species), and those trained from scratch on the target species suggests that the Chicks4FreeID dataset likely has unique characteristics not present in the datasets used to pretrain the MegaDescriptor. Thus, our dataset could enhance the underlying data distribution used to train general animal re-identification models like the MegaDescriptor.

Additionally, there is a small improvement in scores between the Vision Transformer over the Swin architecture, which was used to train the MegaDescriptor. The slightly better performance of the Vision Transformer might be due to two reasons: First, we observed a more stable training process for the Vision Transformer (cross-entropy loss) than for the Swin Transformer (ArcFace loss). Therefore we believe that training a more straightforward approach allows for easier convergence on a small dataset like ours. Second, we replaced the standard classification head of the Vision Transformer with a simple linear layer. Since a simple linear layer has limited discriminative power, achieving good overall performance suggests the presence of good embeddings, which was confirmed by the embedding evaluation using k-NN.

5 Conclusion

5.1 Findings

The Chicks4FreeID benchmark dataset was introduced. To the best of our knowledge, it is the very first publicly available dataset for chicken re-identification. The dataset is well-annotated and released under the relatively unrestrictive CC BY 4.0 license. It contains 1270 instance annotations of 54 individuals - 50 individuals and 1215 of the instances are chicken. The 677 images, which depict mainly chickens from 11 different coops and various breeds, were individually captured rather than derived from video. The dataset was created systematically, with manual annotation and instance-to-individual assignments based on expert knowledge, without the use of automated methods, ensuring reliable ground truth annotations. Instead of providing merely bounding boxes that might include parts of the background or other individuals, we offer preprocessed cut-out crops based on precise segmentations of the instances. While the main use case of the dataset is the re-identification of chickens, it also supports semantic and instance segmentation. In addition to instance and semantic segmentation masks, information on identity, animal category, and coop, the dataset also includes a visibility rating of the instances, accounting for occlusions. For the task of closed set chicken re-identification, we established a baseline on the dataset, achieving Top-1 accuracy scores up to 0.928, Top-5 accuracy scores up to 0.990, and mAP scores up to 0.976 with the Vision Transformer. The experiments suggest that the introduced dataset could be a valuable resource for training more robust (general) animal re-identification systems.

5.2 Limitations

One clear limitation of the dataset is its size. With 1215 instance annotations of 50 chicken individuals, it is comparatively small. There also exists an imbalance within the classes (individuals), with the number of instances ranging from 4 to 27 in the "best" visibility subset. For chicken breeds with minimal inter-individual variability (e.g., uniform plumage), having more individuals and more instances of each individual would likely aid in re-identification. Additionally, all images of a given chicken were taken on the same day, so changes in appearance over time were not captured. An open question is the dataset's applicability to industrial farming, where thousands of chickens of a single breed are typically kept. A specialized dataset for such breeds could potentially be more suitable for commercial applications. Furthermore, the chicken breeds included in the Chicks4FreeID dataset are not exhaustive, despite their variability. The specific breeds were not annotated because they could not always be accurately determined.

5.3 Future work

To further enhance the Chicks4FreeID dataset and address its current limitations, future work could focus on several promising directions. Expanding the dataset to include a larger number of individuals and an even broader range of breeds would enhance its robustness and generalizability. Enriching the metadata with detailed breed-specific information could provide additional context. Methods to automatically create new labeled samples from existing data using generative AI, as proposed in [97], could be evaluated for their potential to aid in expanding the dataset. To capture changes in appearance over time due to factors such as molting, growth, and environmental conditions, individuals from the dataset may be photographed again, provided they are still alive. Similarly, new individuals added to the dataset could be photographed repeatedly over time. The versioning system of the dataset facilitates potential expansions and continuous improvements, ensuring its ongoing relevance and applicability for future research. However, the challenge of long-term data collection persists, as free-range chickens often fall prey to wild predators (e.g., foxes or raccoons). Another interesting direction for future work would be the investigation of models trained on the dataset and their applicability to industrial farming settings with crowded conditions and chickens of a single breed. On a final note, we envision the Chicks4FreeID dataset being utilized by established and aspiring researchers alike, i.e., in future research, contributing to the development of chicken-specific and multi-species re-identification systems, as well as being used for practicing purposes.

Acknowledgments and Disclosure of Funding

We are immensely thankful to the kind chicken owners who opened their coops for our research, allowing us to collect data and generously offering us fresh eggs. Each of your chickens has made a unique and valuable contribution to the advancement of science.

References

[1] A. M. Guhl and L. L. Ortman, "Visual Patterns in the Recognition of Individuals among Chickens," *The Condor*, vol. 55, no. 6, pp. 287–298, 11 1953. [Online]. Available: https://doi.org/10.2307/1365008

[2] E. T. Psota, T. Schmidt, B. Mote, and L. C. Pérez, "Long-term tracking of group-housed livestock using keypoint detection and map estimation for individual animal identification," *Sensors*, vol. 20, no. 13, p. 23, 2020. [Online]. Available: https://doi.org/10.3390/s20133670

[3] W. Lu, Y. Zhao, J. Wang, Z. Zheng, L. Feng, and J. Tang, "Mammalclub: An annotated wild mammal dataset for species recognition, individual identification, and behavior recognition," *Electronics*, vol. 12, no. 21, p. 14, 2023. [Online]. Available: https://www.mdpi.com/2079-9292/12/21/4506

[4] J. Chan, H. Carrión, R. Mégret, J. L. A. Rivera, and T. Giray, "Honeybee re-identification in video: New datasets and impact of self-supervision," in *Proceedings of the 17th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP 2022) - Volume 5: VISAPP*, INSTICC. Avenida de S. Francisco Xavier, Lote 7 Cv. C, 2900-616 Setúbal, Portugal: SciTePress, 2022, pp. 517–525.

[5] M. Vidal, N. Wolf, B. Rosenberg, B. P. Harris, and A. Mathis, "Perspectives on Individual Animal Identification from Biology and Computer Vision," *Integrative and Comparative Biology*, vol. 61, no. 3, pp. 900–916, 05 2021. [Online]. Available: https://doi.org/10.1093/icb/icab107

[6] P. Kulits, J. Wall, A. Bedetti, M. Henley, and S. Beery, "Elephantbook: A semi-automated human-in-the-loop system for elephant re-identification," in *Proceedings of the 4th ACM SIGCAS Conference on Computing and Sustainable Societies*, ser. COMPASS '21. New York, NY, USA: Association for Computing Machinery, 2021, p. 88–98. [Online]. Available: https://doi.org/10.1145/3460112.3471947

[7] O. Moskvyak, F. Maire, F. Dayoub, A. O. Armstrong, and M. Baktashmotlagh, "Robust re-identification of manta rays from natural markings by learning pose invariant embeddings," in *2021 Digital Image Computing: Techniques and Applications (DICTA)*. Piscataway, NJ: IEEE, 2021, pp. 1–8.

[8] L. Wang, R. Ding, Y. Zhai, Q. Zhang, W. Tang, N. Zheng, and G. Hua, "Giant panda identification," *IEEE Transactions on Image Processing*, vol. 30, pp. 2837–2849, 2021.

[9] Q. He, Q. Zhao, N. Liu, P. Chen, Z. Zhang, and R. Hou, "Distinguishing individual red pandas from their faces," in *Pattern Recognition and Computer Vision*, Z. Lin, L. Wang, J. Yang, G. Shi, T. Tan, N. Zheng, X. Chen, and Y. Zhang, Eds. Cham: Springer International Publishing, 2019, pp. 714–724.

[10] W. Andrew, S. Hannuna, N. Campbell, and 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)*. Piscataway, NJ: IEEE, 2016, pp. 484–488.

[11] S. Li, L. Fu, Y. Sun, Y. Mu, L. Chen, J. Li, and H. Gong, "Cow dataset," https://doi.org/10.6084/m9.figshare.16879780, 2021.

[12] W. Andrew, C. Greatwood, and T. Burghardt, "Visual localisation and individual identification of holstein friesian cattle via deep learning," in *2017 IEEE International Conference on Computer Vision Workshops (ICCVW)*. Piscataway, NJ: IEEE, 2017, pp. 2850–2859.

[13] J. Gao, T. Burghardt, W. Andrew, A. W. Dowsey, and N. W. Campbell, "Towards self-supervision for video identification of individual holstein-friesian cattle: The cows2021 dataset," 2021.

[14] W. Andrew, J. Gao, S. Mullan, N. Campbell, A. W. Dowsey, and T. Burghardt, "Visual identification of individual holstein-friesian cattle via deep metric learning," *Computers and Electronics in Agriculture*, vol. 185, p. 106133, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0168169921001514

[15] T. Zhang, Q. Zhao, C. Da, L. Zhou, L. Li, and S. Jiancuo, "Yakreid-103: A benchmark for yak re-identification," in *2021 IEEE International Joint Conference on Biometrics (IJCB)*. Piscataway, NJ: IEEE, 2021, pp. 1–8.

[16] F. Tausch, S. Stock, J. Fricke, and O. Klein, "Bumblebee re-identification dataset," in *2020 IEEE Winter Applications of Computer Vision Workshops (WACVW)*. Piscataway, NJ: IEEE, 2020, pp. 35–37.

[17] P. Borlinghaus, F. Tausch, and L. Rettenberger, "A purely visual re-id approach for bumblebees (bombus terrestris)," *Smart Agricultural Technology*, vol. 3, p. 100135, 2023. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2772375522000995

[18] M. Ye, J. Shen, G. Lin, T. Xiang, L. Shao, and S. C. H. Hoi, “Deep learning for person re-identification: A survey and outlook,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 6, pp. 2872–2893, 2022.

[19] C. L. Witham, “Automated face recognition of rhesus macaques,” *Journal of Neuroscience Methods*, vol. 300, pp. 157–165, 2018, measuring Behaviour 2016. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0165027017302637

[20] A. Freytag, E. Rodner, M. Simon, A. Loos, H. S. Kühl, and J. Denzler, “Chimpanzee faces in the wild: Log-euclidean cnns for predicting identities and attributes of primates,” in *Pattern Recognition*, B. Rosenhahn and B. Andres, Eds. Cham: Springer International Publishing, 2016, pp. 51–63.

[21] T.-Y. Lin and Y.-F. Kuo, “Cat face recognition using deep learning,” St. Joseph, MI, 2018.

[22] N. Dlamini and T. L. v. Zyl, “Lion face dataset,” https://github.com/tvanzyl/wildlife_reidentification/, 2020, mara Masia project, Kenya.

[23] G. Mougeot, D. Li, and S. Jia, “A deep learning approach for dog face verification and recognition,” in *PRICAI 2019: Trends in Artificial Intelligence*, A. C. Nayak and A. Sharma, Eds. Cham: Springer International Publishing, 2019, pp. 418–430.

[24] S. Li, J. Li, H. Tang, R. Qian, and W. Lin, “Atrw: A benchmark for amur tiger re-identification in the wild,” in *Proceedings of the 28th ACM International Conference on Multimedia*, ser. MM ’20. New York, NY, USA: Association for Computing Machinery, 2020, p. 2590–2598. [Online]. Available: https://doi.org/10.1145/3394171.3413569

[25] J. Parham, J. Crall, C. Stewart, T. Berger-Wolf, and D. Rubenstein, “Animal population censusing at scale with citizen science and photographic identification,” *AAAI Spring Symposium - Technical Report*, vol. SS-17-01 - SS-17-08, pp. 37 – 44, 2017. [Online]. Available: http://arks.princeton.edu/ark:/88435/pr1s791

[26] J. B. Haurum, A. Karpova, M. Pedersen, S. H. Bengtson, and T. B. Moeslund, “Re-identification of zebrafish using metric learning,” in *2020 IEEE Winter Applications of Computer Vision Workshops (WACVW)*. Piscataway, NJ: IEEE, 2020, pp. 1–11.

[27] M. Lahiri, C. Tantipathananandh, R. Warungu, D. I. Rubenstein, and 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*, ser. ICMR ’11. New York, NY, USA: Association for Computing Machinery, 2011. [Online]. Available: https://doi.org/10.1145/1991996.1992002

[28] L. Adam, V. Čermák, K. Papafitsoros, and L. Picek, “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)*. Piscataway, NJ: IEEE, January 2024, pp. 7146–7156.

[29] S. Bouma, M. D. Pawley, K. Hupman, and A. Gilman, “Individual common dolphin identification via metric embedding learning,” in *2018 International Conference on Image and Vision Computing New Zealand (IVCNZ)*. Piscataway, NJ: IEEE, 2018, pp. 1–6.

[30] B. Hughes and T. Burghardt, “Automated visual fin identification of individual great white sharks,” *International Journal of Computer Vision*, vol. 122, no. 3, pp. 542–557, May 2017. [Online]. Available: https://doi.org/10.1007/s11263-016-0961-y

[31] H. B. Bae, D. Pak, and S. Lee, “Dog nose-print identification using deep neural networks,” *IEEE Access*, vol. 9, pp. 49 141–49 153, 2021.

[32] M. Zuerl, R. Dirauf, F. Koeferl, N. Steinlein, J. Sueskind, D. Zanca, I. Brehm, L. v. Fersen, and B. Eskofier, “Polarbearvidid: A video-based re-identification benchmark dataset for polar bears,” *Animals*, vol. 13, no. 5, p. 16, 2023. [Online]. Available: https://www.mdpi.com/2076-2615/13/5/801

[33] M. Clapham, E. Miller, M. Nguyen, and C. T. Darimont, “Automated facial recognition for wildlife that lack unique markings: A deep learning approach for brown bears,” *Ecology and Evolution*, vol. 10, no. 23, pp. 12 883–12 892, 2020. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/ece3.6840

[34] M. Körschens and J. Denzler, “Elpephants: A fine-grained dataset for elephant re-identification,” in *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*. Piscataway, NJ: IEEE, 2019, pp. 263–270.

[35] V. Čermák, L. Picek, L. Adam, and 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)*. Piscataway, NJ: IEEE, January 2024, pp. 5953–5963.

[36] O. Wahltinez and 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*, vol. 15, no. 2, pp. 373–387, 2024.

[37] S. Guo, P. Xu, Q. Miao, G. Shao, C. A. Chapman, X. Chen, G. He, D. Fang, H. Zhang, Y. Sun, Z. Shi, and B. Li, “Automatic identification of individual primates with deep learning techniques,” *iScience*, vol. 23, no. 8, p. 32, Aug 2020. [Online]. Available: https://doi.org/10.1016/j.isci.2020.101412

[38] A. C. Ferreira, L. R. Silva, F. Renna, H. B. Brandl, J. P. Renoult, D. R. Farine, R. Covas, and C. Doutrelant, “Bird individualid,” https://github.com/AndreCFerreira/Bird_individualID, 2020.

[39] L. I. Kuncheva, F. Williams, S. L. Hennessey, and J. J. Rodríguez, “Animal-identification-from-video,” https://github.com/LucyKuncheva/Animal-Identification-from-Video, 2022.

[40] L. Adam, V. Čermák, K. Papafitsoros, and L. Picek, “Seaturtleid,” https://www.kaggle.com/datasets/wildlifedatasets/seaturtleid2022, 2022.

[41] O. Wahltinez, “Sea star re-id,” https://lila.science/sea-star-re-id-2023/, 2023.

[42] W. Me, “Beluga id,” https://lila.science/datasets/beluga-id-2022/, 2022, info@wildme.org.

[43] T. Cheeseman, K. Southerland, W. Reade, and A. Howard, “Happywhale - whale and dolphin identification,” https://kaggle.com/competitions/happy-whale-and-dolphin, 2022.

[44] E. Nepovinnykh, “Sealid,” https://doi.org/10.23729/0f4a3296-3b10-40c8-9ad3-0cf00a5a4a53, 2022, lappeenranta University of Technology, School of Engineering Science yhteiset.

[45] K. Papafitsoros, L. Adam, V. Čermák, and L. Picek, “Seaturtleid,” https://www.kaggle.com/datasets/wildlifedatasets/seaturtleidheads, 2022.

[46] W. T. Watch and L. O. Conservation, “Turtle recall: Conservation challenge,” https://zindi.africa/competitions/turtle-recall-conservation-challenge/data, 2022.

[47] C. Trotter, G. Atkinson, M. Sharpe, K. Richardson, A. S. McGough, N. Wright, B. Burville, and P. Berggren, “The northumberland dolphin dataset 2020,” https://doi.org/10.25405/data.ncl.c.4982342, 2020, newcastle Uniyversity. Collection.

[48] A. Howard, inversion, K. Southerland, and T. Cheeseman, “Humpback whale identification,” 2018. [Online]. Available: https://kaggle.com/competitions/humpback-whale-identification

[49] C. B. Khan, Shashank, and W. Kan, “Right whale recognition,” 2015. [Online]. Available: https://kaggle.com/competitions/noaa-right-whale-recognition

[50] J. Holmberg, B. Norman, and Z. Arzoumanian, “Whale shark id,” https://lila.science/datasets/whale-shark-id, 2020, info@wildme.org.

[51] J. Gao, T. Burghardt, W. Andrew, A. W. Dowsey, and N. W. Campbell, “Cows2021,” https://doi.org/10.5523/bris.4vnrca7qw1642qlwxjadp87h7, 2021.

[52] W. Andrew, J. Gao, S. Mullan, N. Campbell, A. W. Dowsey, and T. Burghardt, “Opencows2020,” https://doi.org/10.5523/bris.10m32xl88x2b61zlkkgz3fml17, 2020.

[53] W. Andrew, C. Greatwood, and T. Burghardt, “Aerialcattle2017,” 2017. [Online]. Available: https://doi.org/10.5523/bris.3owflku95bxsx24643cybxu3qh

[54] ——, “Friesiancattle2017,” 2017. [Online]. Available: https://doi.org/10.5523/bris.2yizcfbkuv4352pzc32n54371r

[55] W. Andrew, S. Hannuna, N. Campbell, and T. Burghardt, “Friesiancattle2015,” 2016. [Online]. Available: https://doi.org/10.5523/bris.wurzq71kfm561ljahbwjhx9n3

[56] D. Kern, T. Schiele, U. Klauck, A. DeSilva, and W. Ingabire, “Chicks4freeid,” https://huggingface.co/datasets/dariakern/Chicks4FreeID, 2024.

[57] W. Lu, Y. Zhao, J. Wang, Z. Zheng, L. Feng, and J. Tang, “Mammalclub,” https://github.com/WJ-0425/MammalClub, 2023, download: https://pan.baidu.com/s/1in8xJxdjoNMNb3yuKbVDPA. (code: o5tq).

[58] Z. He, J. Qian, D. Yan, C. Wang, and Y. Xin, “Animal re-identification algorithm for posture diversity,” in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Piscataway, NJ: IEEE, 2023, pp. 1–5.

[59] Z. He, "Multi-pose dog dataset," https://data.mendeley.com/datasets/v5j6m8dzhv/1, 2023, mendeley Data, V1.

[60] M. Zuerl, R. Dirauf, F. Koeferl, N. Steinlein, J. Sueskind, D. Zanca, I. Brehm, L. von Fersen, and B. Eskofier, "PolarBearVidID: A Video-based Re-Identification Benchmark Dataset for Polar Bears," Jan. 2023. [Online]. Available: https://doi.org/10.5281/zenodo.7564529

[61] L. I. Kuncheva, F. Williams, S. L. Hennessey, and J. J. Rodríguez, "A benchmark database for animal re-identification and tracking," in *2022 IEEE 5th International Conference on Image Processing Applications and Systems (IPAS)*, vol. Five. Piscataway, NJ: IEEE, 2022, pp. 1–6.

[62] B. P. C. Trust, "Hyiena id," https://lila.science/datasets/hyena-id-2022/, 2022, panthera pardus CSV custom export. Retrieved from African Carnivore Wildbook 2022-04-28. info@wildme.org.

[63] ——, "Leopard id," https://lila.science/datasets/leopard-id-2022/, 2022, panthera pardus CSV custom export. Retrieved from African Carnivore Wildbook 2022-04-28. info@wildme.org.

[64] E. Nepovinnykh, T. Eerola, V. Biard, P. Mutka, M. Niemi, M. Kunnasranta, and H. Kälviäinen, "Sealid: Saimaa ringed seal re-identification dataset," *Sensors*, vol. 22, no. 19, p. 11p, 2022. [Online]. Available: https://www.mdpi.com/1424-8220/22/19/7602

[65] K. Papafitsoros, L. Adam, V. Čermák, and L. Picek, "Seaturtleid: A novel long-span dataset highlighting the importance of timestamps in wildlife re-identification," 2022.

[66] S. Li, L. Fu, Y. Sun, Y. Mu, L. Chen, J. Li, and H. Gong, "Individual dairy cow identification based on lightweight convolutional neural network," *PLoS ONE*, vol. 16, no. 11, p. 13, 2021.

[67] V. Miele, G. Dussert, B. Spataro, S. Chamaillé-Jammes, D. Allainé, and C. Bonenfant, "Revisiting animal photo-identification using deep metric learning and network analysis," *Methods in Ecology and Evolution*, vol. 12, no. 5, pp. 863–873, 2021. [Online]. Available: https://besjournals.onlinelibrary.wiley.com/doi/abs/10.1111/2041-210X.13577

[68] ——, "Giraffe dataset," https://plmlab.math.cnrs.fr/vmiele/animal-reid/, 2020, ftp link ftp://pbil.univ-lyon1.fr/pub/datasets/miele2021.

[69] L. Wang, R. Ding, Y. Zhai, Q. Zhang, W. Tang, N. Zheng, and G. Hua, "ipanda-50," https://github.com/iPandaDateset/iPanda-50, 2021.

[70] J. B. Haurum, A. Karpova, M. Pedersen, S. H. Bengtson, and T. B. Moeslund, "Aau zebrafish re-identification dataset," https://www.kaggle.com/datasets/aalborguniversity/aau-zebrafish-reid, 2020.

[71] S. Guo, P. Xu, Q. Miao, G. Shao, C. A. Chapman, X. Chen, G. He, D. Fang, H. Zhang, Y. Sun, Z. Shi, and B. Li, "Afd," https://doi.org/10.17632/z3x59pv4bz.2, 2020, mendeley Data Version 2.

[72] S. Li, J. Li, H. Tang, R. Qian, and W. Lin, "Atrw (amur tiger re-identification in the wild)," https://lila.science/datasets/atrw, 2020.

[73] N. Dlamini and T. L. v. Zyl, "Automated identification of individuals in wildlife population using siamese neural networks," in *2020 7th International Conference on Soft Computing and Machine Intelligence (ISCMI)*. Piscataway, NJ: IEEE, 2020, pp. 224–228.

[74] C. Trotter, G. Atkinson, M. Sharpe, K. Richardson, A. S. McGough, N. Wright, B. Burville, and P. Berggren, "NDD20: A large-scale few-shot dolphin dataset for coarse and fine-grained categorisation," *CoRR*, vol. abs/2005.13359, p. 5, 2020. [Online]. Available: https://arxiv.org/abs/2005.13359

[75] N. Dlamini and T. L. v. Zyl, "Nyala dataset," https://github.com/tvanzyl/wildlife_reidentification/, 2020, south African nature reserves.

[76] A. C. Ferreira, L. R. Silva, F. Renna, H. B. Brandl, J. P. Renoult, D. R. Farine, R. Covas, and C. Doutrelant, "Deep learning-based methods for individual recognition in small birds," *Methods in Ecology and Evolution*, vol. 11, no. 9, pp. 1072–1085, 2020. [Online]. Available: https://besjournals.onlinelibrary.wiley.com/doi/abs/10.1111/2041-210X.13436

[77] G. Mougeot, D. Li, and S. Jia, "Dog face dataset," https://github.com/GuillaumeMougeot/DogFaceNet, 2019.

[78] T.-Y. Lin and Y.-F. Kuo, "Cat individual images," https://www.kaggle.com/datasets/timost1234/cat-individuals, 2018.

[79] J. Schneider, N. Murali, G. Taylor, and J. Levine, "Can drosophila melanogaster tell who's who?" *PLoS ONE*, vol. 13, no. 10, p. 10, 2018.

[80] ——, "Dataset for: Can Drosophila melanogaster tell who's who?" 2018. [Online]. Available: https://doi.org/10.5683/SP2/JP4WDF

[81] C. L. Witham, "Macaquefaces," https://github.com/clwitham/MacaqueFaces, 2018.

[82] J. Parham, J. Crall, C. Stewart, T. Berger-Wolf, and D. Rubenstein, "Great zebra and giraffe count id," https://lila.science/datasets/great-zebra-giraffe-id, 2017, info@wildme.org.

[83] A. Freytag, E. Rodner, M. Simon, A. Loos, H. S. Kühl, and J. Denzler, "Chimpanzee faces in the wild," https://github.com/cvjena/chimpanzee_faces, 2016, acknowledgements: Tobias Deschner, Laura Aporius, Karin Bahrke, Zoo Leipzig.

[84] J. Holmberg, B. Norman, and Z. Arzoumanian, "Estimating population size, structure, and residency time for whale sharks rhincodon typus through collaborative photo-identification," *Endangered Species Research*, vol. 7, no. 1, pp. 39–53, 2009.

[85] Labelbox, ""labelbox"," [Online] Available: https://labelbox.com, 2024.

[86] E. D. Cubuk, B. Zoph, J. Shlens, and Q. Le, "Randaugment: Practical automated data augmentation with a reduced search space," in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020, pp. 18 613–18 624. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2020/file/d85b63ef0ccb114d0a3bb7b7d808028f-Paper.pdf

[87] I. Susmelj, M. Heller, P. Wirth, J. Prescott, and M. E. et al., "Lightly," 2020.

[88] V. Čermák, L. Picek, L. Adam, and K. Papafitsoros, "Megadescriptor-l-384," https://huggingface.co/BVRA/MegaDescriptor-L-384, 2024.

[89] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, "Swin transformer: Hierarchical vision transformer using shifted windows," in *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*. Los Alamitos, CA, USA: IEEE Computer Society, oct 2021, pp. 9992–10 002. [Online]. Available: https://doi.ieeecomputersociety.org/10.1109/ICCV48922.2021.00986

[90] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "Arcface: Additive angular margin loss for deep face recognition," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.

[91] M. Oquab, T. Darcet, T. Moutakanni, H. Vo, M. Szafraniec, V. Khalidov, P. Fernandez, D. Haziza, F. Massa, A. El-Nouby, M. Assran, N. Ballas, W. Galuba, R. Howes, P.-Y. Huang, S.-W. Li, I. Misra, M. Rabbat, V. Sharma, G. Synnaeve, H. Xu, H. Jegou, J. Mairal, P. Labatut, A. Joulin, and P. Bojanowski, "Dinov2: Learning robust visual features without supervision," 2024, unpublished.

[92] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever, "Learning transferable visual models from natural language supervision," 2021, unpublished.

[93] R. Wightman, "Pytorch image models," https://github.com/rwightman/pytorch-image-models, 2019.

[94] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An image is worth 16x16 words: Transformers for image recognition at scale," in *International Conference on Learning Representations*, 2021. [Online]. Available: https://openreview.net/forum?id=YicbFdNTTy

[95] T. maintainers and contributors, "Torchvision: Pytorch's computer vision library," https://github.com/pytorch/vision, 2016.

[96] N. S. Detlefsen, J. Borovec, J. Schock, A. H. Jha, T. Koker, L. D. Liello, D. Stancl, C. Quan, M. Grechkin, and W. Falcon, "Torchmetrics - measuring reproducibility in pytorch," *Journal of Open Source Software*, vol. 7, no. 70, p. 4101, 2022. [Online]. Available: https://doi.org/10.21105/joss.04101

[97] Y. Zhang, D. Zhou, B. Hooi, K. Wang, and J. Feng, "Expanding small-scale datasets with guided imagination," in *Advances in Neural Information Processing Systems*, A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, Eds., vol. 36. Curran Associates, Inc., 2023, pp. 76 558–76 618. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2023/file/f188a55392d3a7509b0b27f8d24364bb-Paper-Conference.pdf

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] See Section 5.2.
   (c) Did you discuss any potential negative societal impacts of your work? [N/A] It is a chicken dataset.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments (e.g. for benchmarks)...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] See supplementary material.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Section 4.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] See Table 2.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See Section 4.2.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] We used existing architectures, models and code. The creators were cited.
   (b) Did you mention the license of the assets? [Yes] We mentioned the licenses in the paper and in the supplementary material.
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] The Chicks4FreeID dataset. https://doi.org/10.57967/hf/2345
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]