

# CNN Based Flank Predictor for Quadruped Animal Species

Vanessa Suessle<sup>1,\*</sup> Marco Heurich<sup>2,3,4</sup> Colleen T. Downs<sup>5</sup> Andreas Weinmann<sup>6</sup> Elke Hergenroether<sup>1</sup>

<sup>1</sup> Department of Computer Science, University of Applied Sciences Darmstadt, Schoefferstrasse 3, Darmstadt, Germany

<sup>2</sup> Department of National Park Monitoring and Animal Management, Bavarian Forest National Park, Freyunger Str. 2, 94481 Grafenau, Germany

<sup>3</sup> Department of Wildlife Ecology and Management, Faculty of Environment and Natural Resources, University of Freiburg, Freiburg, Germany

<sup>4</sup> Faculty of Applied Ecology, Agricultural Sciences and Biotechnology, Inland Norway University of Applied Sciences, Evenstad, Norway

<sup>5</sup> School of Life Sciences, University of KwaZulu-Natal, Carbis Road, Scottsville, Pietermaritzburg, South Africa

<sup>6</sup> Department of Mathematics, University of Applied Sciences Darmstadt, Schoefferstrasse 3, Darmstadt, Germany

\* E-mail: vanessa.suessle@h-da.de


**Abstract:** The bilateral asymmetry of flanks of animals with visual body marks that uniquely identify an individual, complicates tasks like population estimations. Automatically generated additional information on the visible side of the animal would improve the accuracy for individual identification. In this study we used transfer learning on popular CNN image classification architectures to train a flank predictor that predicts the visible flank of quadruped mammalian species in images. We automatically derived the data labels from existing datasets originally labeled for animal pose estimation. We trained the models in two phases with different degrees of retraining. The developed models were evaluated in different scenarios of different unknown quadruped species in known and unknown environments. As a real-world scenario, we used a dataset of manually labeled Eurasian lynx (*Lynx lynx*) from camera traps in the Bavarian Forest National Park to evaluate the model. The best model, trained on an EfficientNetV2 backbone, achieved an accuracy of 88.70 % for the unknown species lynx in a complex habitat.

## 1 Introduction

In recent years, camera traps have emerged as an important tool for wildlife monitoring in the context of wildlife conservation and management. They are low cost, non-invasive, capture the animal in its natural habitat and can be triggered at any time during night and day. However, when multiple camera traps are deployed, the amount of captured data quickly surpasses what can be manually analyzed. Analyzing such data is time-intensive and prone to errors, especially when conducted manually over several hours [1]. In addition to the common tasks of detection and species identification, researchers can gain additional information from the captured data, such as the visible side of the animal. The additional information can improve main analysis tasks like individual identification, which is of special interest to monitor the population of a particular species in a specific area. It relies on visual characteristics that uniquely identify an individual, for example fur patterns for lynx, zebras (*Equus quagga*) and leopards (*Panthera pardus*) or body marks on manta rays (*Manta alfredi*) and whale sharks (*Rhincodon typus*) [2]. However, a challenge in individual identification arises because the visual characteristics are usually bilaterally asymmetric on the flanks of an animal [3]. Meaning they are independent and both sides of an individual need to be captured for a successful individual identification. If solely each side is captured individually, with no overlap in videos or images, the flanks cannot accurately be assigned to the same individual. To cope with the issue of the bilateral asymmetric flanks, some studies only concentrate on one flank of all individuals for the analysis, not using the full potential of the data [3, 4].

Other studies treat both sides independently as individuals [5], but it bears the risk of overestimating the population, which can bias the results of a monitoring and therewith can negatively affect management decisions. Alternatively statistical methods can be applied for approximations with missing flank footage [6].

Over the past decades, many studies have mentioned the challenge of bilaterally asymmetric individual markings for individual identification [3, 4, 8]. To address this challenge studies on Felidae monitoring often use a dual camera setup from opposing positions [9]. Even though this approach manages to capture both flanks of an animal, the subsequent manual analysis required to sort the captions



**Fig. 1:** Example labels of body keypoints for pose estimation [7]

into the classes on the visible flank remains time-intensive. Therefore, incorporating a component for predicting the visible flank of an animal could save labeling time, prevent overestimating of populations and could help to get the full potential from the captured data. A flank predictor could be embedded in an individual identification pipeline as part of the data preprocessing to increase accuracy by filtering out impossible matches on opposite flanks, reduce computing time and store additional information on the available viewpoints on individuals, because for research on population monitoring not only matching two images as one individual is relevant, but also to clearly identify to images as impossible to match [10]. For existing datasets collected with a single camera at a specific location, it is not possible to retrospectively obtain the information that could have been extracted from a setup with two opposing cameras. Furthermore, future studies with budget limitations and a restricted number of camera traps may have to choose between a larger monitoring area or having dual captions at a camera spot. Apart from the concept of dual camera trap setup, no automatization exists for identifying the visible flank, which could significantly enhance wildlife monitoring in terms of analysis time and accuracy. This study aims to develop a model that can be applied to multiple species and integrated as a modular component in individual identification pipelines. Training a convolutional neural network (CNN) model for flank prediction requires a large amount of labeled data, which is time-consuming to produce manually. In this work an approach is used to automatically derive flank labels from available datasets (Table 1) of quadruped mammals that were labelled with skeletal keypoints (Figure 1). Based on the assumption that even though animals have a high diversity in their appearances, the basic skeletal structure and locomotion of quadrupeds is similar.

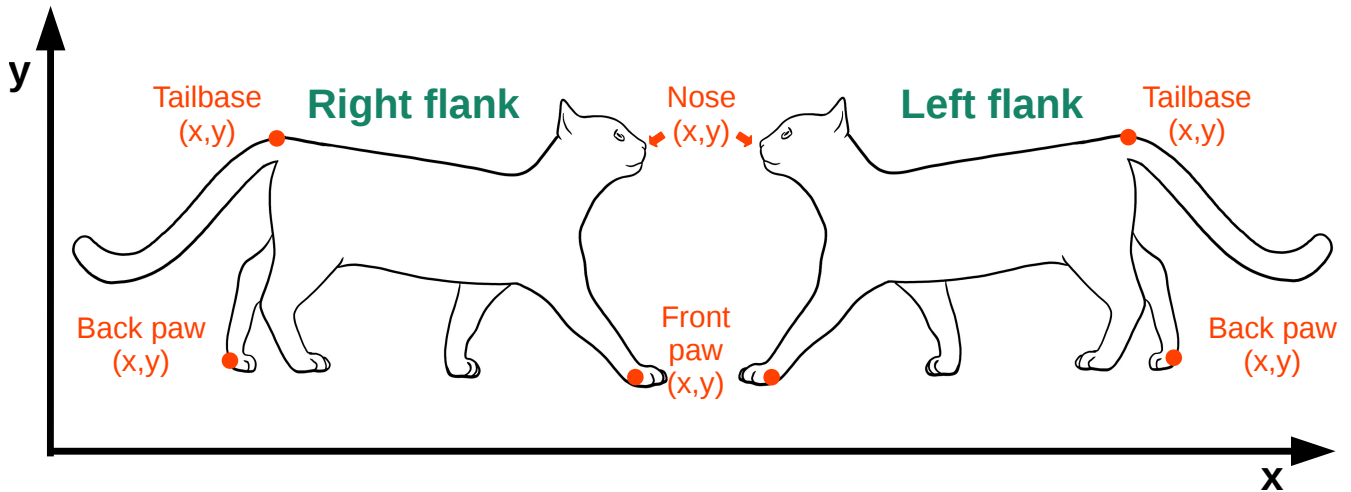


Fig. 2: Flank label based on the relative horizontal position of the body keypoints.

## 2 Related Work

Many studies focused on the systematic setup of camera traps to find suitable locations to increase the possibility to capture usable images. Favourable camera spots may be close to water sources, on trails, ridges or locations that were known to be frequently used by the species of interest from former observations or data from collared individuals [3]. The ideal locations highly depend on the topographical aspects and species of interest and should be determined in consultation with local experts [11]. To increase the chances of capturing camera spots can be equipped with baits to lure wildlife. A dual setup on each side of the trail increases the chance of capturing both flanks of an individual [3, 8, 9]. For single flank captions it is of interest to define which flank is visible for further analysis. But, even for datasets that were collected in a dual camera trap setup, the manual task to label one of the images of each image pair on their viewpoint, to automatically label the image from the opposite side remains. To benefit the actual analysis objective of the research, which could focus on individual identification or behavioural studies, several studies added interim analysis steps during the data processing that would return additional insight into the data. Those additional tasks that many researchers have in common may include the detection [12] and localization of animals in images, counting individuals in herds [13] or pose estimation or tracking on the seen animal [7]. All those approaches are related to the flank prediction focused on in this work and share the challenge of high diversity in the nature of animals. Many of the solutions are bound to a species or a group of species. Pre-trained models mostly exist for domestic or farm animals because of the easier collection of data.

For the task of animal detection a generic model for cross-species application with a high accuracy in detection is the MegaDetector. Its high adaptability to versatile habitats and applications makes it an ideal first component for many analysis pipelines [14]. A similarly versatile model for multi species applications for a generic flank prediction would benefit a multitude of research tasks. Pose estimation is a closely related task to flank prediction. It predicts the pose of an animal shown in a video or image based on predefined body keypoints such as joints or other prominent characteristics like eyes and nose, similar to human pose estimation [15]. Most approaches are tied to the species the model was trained on. The first model that is applicable to versatile cross-species applications is DeepLab's SuperAnimal-Quadruped [16] model which has been trained on over 40,000 images of four-legged animals. It was planned as an out of the box model to suit many applications without data labelling. A binary classification flank prediction seems less complex than an exact pose estimation on several keypoints, and can be held more generic for cross-species applications. Many pose estimation models rely on video data inputs. Due to higher battery consumption and

storage image data is often preferred over video data when collecting data in the wild. Although, meta data from consecutive frames in unlabeled video data can benefit individual identification compared to image data [10].

To the best of our knowledge the only study that tackled the issue of predicting the captured flank of an animal, developed a model that classifies images into different classes combining the species and viewpoint. The class included four different animal species and eight perspectives for zebras and giraffes (*Giraffa reticulata*). The model was trained on combined species-viewpoint classes. Examples of the classes are 'zebra-right', 'zebra-left', 'zebra-front', 'giraffe-left' and 'giraffe-front' with a total of 16 combinations for the species zebra and giraffe. Further treated species in the study were whales and turtles [17]. Combining a species classifier and viewpoint prediction has the downside of limiting the application to the species the model was trained on. The obstacle so far was that training a CNN model requires a large amount of labelled training data, which is not available for the classification of the visible flank of other animal species in images. But several datasets labeled with keypoints on poses for different animal species already exist and could be used to extract labels. The available labeled datasets cover a similar set of body keypoints, including joints and prominent features, and cover the following species: different domestic and farm animals [18], different dog breeds [19], horses (*Equus ferus caballus*) [20], amur tigers (*Panthera tigris tigris*) [5], cheetahs (*Acinonyx jubatus*) [7] and a dataset of 50 different mammalian species [21].

Table 1 Datasets with labeled data

Dataset (Species)	Labeled Annotations	Distribution Left / Right Undefined
ATRW [5] (Amur tiger)	2192	1166 / 995 / 32
Animal pose dataset [18] (Dog, cat, sheep, cow, horse)	5660	1973 / 2144 / 1543
AP-10K [21] (50 mammalian species)	8524	2730 / 2844 / 2950
Stanford dogs [19] (Dog breeds)	10,059	4482 / 3936 / 1641

Table 2 Body keypoints sorted into Front and Back class

Front of Body	Back of Body
Ears, nose, head, whiskers, chin, throat, neck, front paws, elbows, shoulder, withers	Tailbase, tailend, back paws, back knees, hip

### 3 Methodology

#### 3.1 Extraction of Data Labels

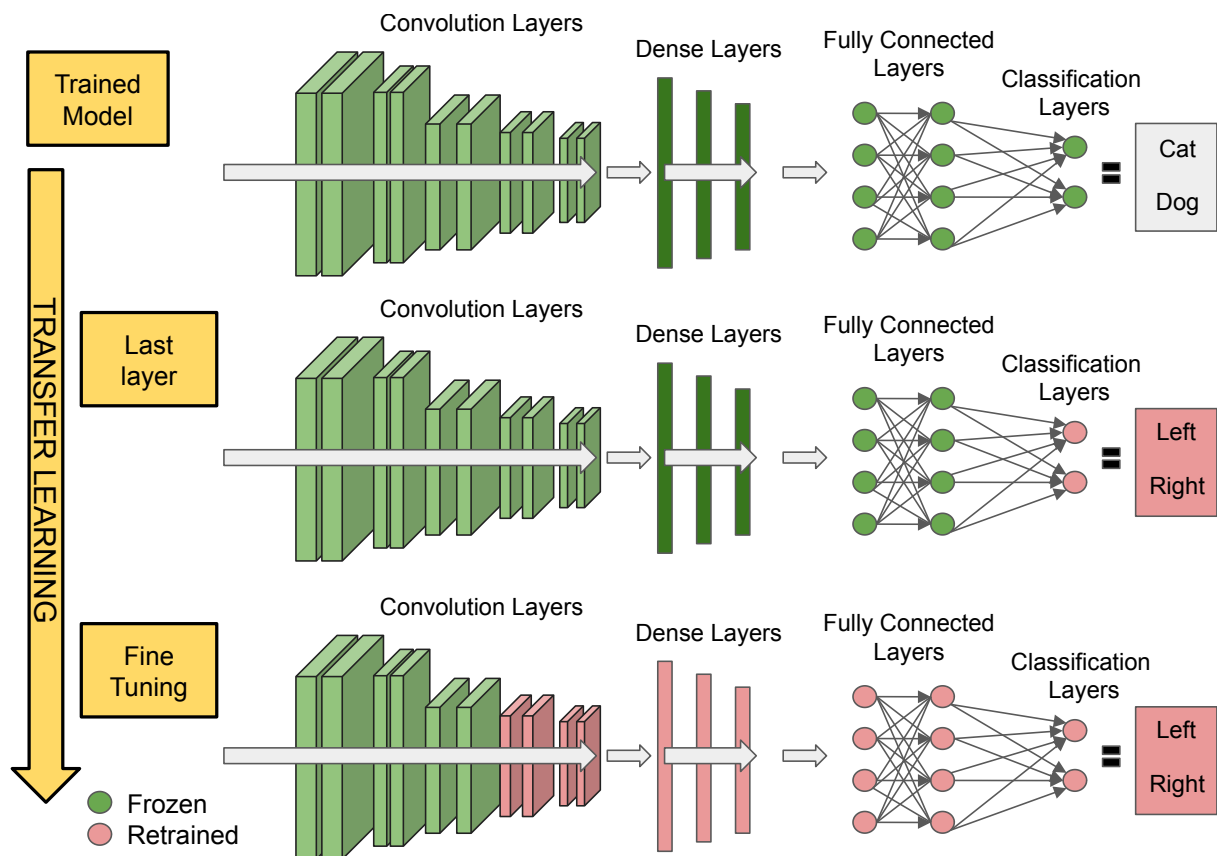
We extracted the labels for the data to train the flank predictor model from different available datasets (Table 1) that had labels for the location of predefined body keypoints (Figure 1). The body keypoints from the different datasets varied, but all datasets had labels for prominent features in common on the head, face and tail as well as on the front and back limbs in common that all quadruped mammalian species share. To obtain a large versatile dataset with flank labels to train a generic model that can be applied to different species, we used multiple datasets with its body keypoints from over 50 different animal species and different dog breeds and combined them in one dataset. The total dataset held 26,000 images showing over 31,000 animals. We derived the labels for each image from the body keypoints, based on the assumption that the animals have a quadruped gait, with a mostly walking, not climbing, locomotion. We extracted the flank labels based on the relative position of the relevant body keypoints in the image. The most important body keypoints were the head and tail, but we considered the position of the front and back paws as well, in case the head and tail were not visible. The keypoint classes from the original datasets were split into two groups. One group included the keypoints that were predefined to belong to the front of the animal's body and the other group of keypoints belonged to the back part of the animal. Table 2 shows example bodypoints for each class.

Depending on the position of the x-value of the points we automatically assigned labels in the following way. If the nose and/or front paws of the animal were located farther on the right side of the image than the tailbase and/or back paws, we labeled the image as 'right', meaning that the right flank of the animal in the image is visible. We proceed analogously for the left side (Figure 2). Images for

which the front and back class keypoints could not clearly be separated across the x-axis and therefore did not fulfill the condition, were not considered for the further process. Table 1 gives a quantitative overview of labels that were successfully extracted from each of the used dataset. The largest proportion of the dataset consists of dogs. We excluded some species from the AP-10K dataset, because they did not fulfill the requirement of pure quadrupedalism or extensively use climbing for locomotion. For an animal hanging vertically on a tree the label cannot clearly be derived based on the x-value position of body front and body back parts (Figure 9). Furthermore, species that were mostly covered by water were filtered out, because of the invisibility of multiple bodyparts. Excluded species were: hippos (*Hippopotamus amphibius*), otters (*Lutrinae*) and any kind of monkeys (*Simiiformes*).

#### 3.2 Training the Model

To develop the flank predictor, we conducted transfer learning on common image classification models. We used ResNet-50 [22], MobileNetV2 [23] and EfficientNetV2-S [24] as backbones, because they are known benchmark models for image classification. All training and validation datasets were preprocessed with the MegaDetector [14] to detect and cut out the annotation of the animal, if bounding boxes on the animal were not already part of the labels. If images contained multiple animals, we processed each annotation individually with an individual flank label. Furthermore, we applied data augmentation techniques to prevent overfitting and to generalize the model. We used random zoom and small random rotations (Figure 5). For the random zoom we either zoomed into the image or out by adding some pixels around the image. For our application case the random rotation was limited to a maximum of 90 degrees, so that the correct flank label and assumptions still apply. The rotation of images strengthens the model's performance and simulates




**Fig. 3:** CNN architecture for transfer learning.

Top: Original architecture.

Middle: All layers frozen, but last layer.

Bottom: Only the first layers are frozen and the later layers are retrained.



**Fig. 4:** Dual camera trap setup to capture both flanks of an individual of a lynx.


real-world cases, because of movement animals will not always be oriented perfectly in a caption. We did not use the data augmentation technique of flipping images, because the mirroring around an axis would lead to an incorrect label.

The idea of transfer learning is to use the weights and layers of existing models and retrain the model on newly labeled data. During this process different layers of the existing models can be frozen, meaning that the parameters in those layers will not be updated. The first layers of a model extract basic features from the processed images. The layers in the rearmost of the network are more specialized than the ones in the front and learn more complex features. The last fully connected layer fulfills the task of classification. It returns a "probability" for each of the classes of the processed image belonging to that class, from which the maximum is assigned as the predicted class. In this study the last layer is a binary classification layer, sorting into the classes 'right' and 'left' stating which flank of the animal is visible in the image. This study combined two approaches in two training phases. The first training phase kept most of the models' layers untouched and only added the last layer for the binary classification. The second phase of training allowed more layers to be retrained by freezing the first half of the layers of each model and adding the fully connected binary classification layer as the last layer. The schematic outline of the two training phases is illustrated in Figure 3.

During the first training phase all layers, but the last layer were frozen. In this phase we trained the models in 15 epochs with a batch size of 32. We added a fully connected dense layer as a last layer to the model backbones. The last binary classification layer was trained on the new labeled data and returns the probability of the processed image belonging to each of the two classes 'left' and 'right'. In the second training step we fine-tuned the model and the weights of multiple later layers in the model were unfrozen and retrained. During retraining, the last layers are fine-tuned and learn features that are more related to the dataset at hand, compared to the earlier learned more generic features. During the fine-tuning phase the first half of layers of the backbones remain untouched to retain their ability to extract basic features. In the second training phase we trained models for further 15 epochs with a batch size of 32 as well. Figure 7 shows the training process over the epochs. The step between the first and second training phase is marked with a red line.

## 4 Case Study

We evaluated the model in two different scenarios on unknown datasets. The two different scenarios covered different degrees of complexity in terms of lighting, quality and poses of the animals. The species present in the two validation dataset were excluded from the training set. For the species unknown to the model we selected Felidae for this research, because many studies focus on the individual identification of Felidae due to declining populations [25].




**Fig. 5:** Data augmentation on the training dataset. The annotations of the cheetah were randomly zoomed in and randomly rotated by less than 90 degrees.[21]

### 4.1 Leopard and Bobcat Dataset

The first validation dataset consists of leopards and bobcats (*Lynx rufus*) and the second datasets covers Eurasian lynx (Figure 6). We split the leopard and bobcat validation dataset from the AP-10K dataset [21], which covers over 50 species. The animals are centered in the images with good viewpoints. The species leopard and bobcat were exclusive in the validation set and not presented to the model during the training process. The leopard and bobcat validation datasets covered similar conditions in terms of quality and lighting as the training dataset.

### 4.2 Lynx Dataset from Cameratraps

The second dataset was collected in the Bavarian Forest National Park and filtered for lynx with two camera traps positioned on opposite sides of a trail to capture individuals from both sides (Figure 4) [26]. The dataset were manually labeled for the viewpoint on the animal and preprocessed with the MegaDetector to detect the animal returning a bounding box. The dataset consists of 4134 images with 2045 being in the class left flank and 2089 in right flank.



**Fig. 6:** Predicted flank labels on the validation datasets for the unseen species: lynx, bobcat and leopard. The prediction highlighted in red is incorrect. Image reference: Bobcat and leopard: [21]; Eurasian lynx: Bavarian Forest National Park

**Table 3** Accuracies obtained by the models on the datasets

Model	Total number / convolutional layers	Frozen layers	Train Accuracy	Validation Accuracy	
				On Unseen Species Leopard / Bobcat	On Unseen Species <b>Lynx</b>
Before fine tuning (one trainable layer)					
MobileNetV2 [23]	35	1	53.29 %	48.87 %	<b>63.09 %</b>
ResNet-50 [22]	50	1	56.77 %	57.89 %	<b>39.01 %</b>
EfficientNetV2-S [24]	42	1	52.72 %	44.36 %	<b>44.65 %</b>
<b>Final models</b> after fine-tuning (half of model's layers frozen)					
MobileNetV2	35	18	90.02 %	97.74 %	<b>71.63 %</b>
ResNet-50	50	23	92.79 %	97.74 %	<b>77.59 %</b>
EfficientNetV2-S	42	20	96.34 %	98.50 %	<b>88.70 %</b>

## 5 Results

We retrained and fine-tuned three models based on the ResNet-50, MobileNetV2 and EfficientNetV2-S backbones using the derived flank labels from datasets including different mammalian species. The models output a prediction on a binary classification task on the visible flank. Exemplary presentations of predictions on images from the lynx dataset and the leopard and bobcat dataset are shown in Figure 6.

### 5.1 Data Augmentation

The graph in Figure 7 shows the training process exemplary for the ResNet-50 with and without data augmentation validated on the lynx dataset. The data augmentation increased the performance of the ResNet-50. Although, the training accuracy without data augmentation was higher than for the model trained with augmented images, the model trained with augmentation performed better on the validation dataset, preventing the model from overfitting.

### 5.2 Fine-Tuning

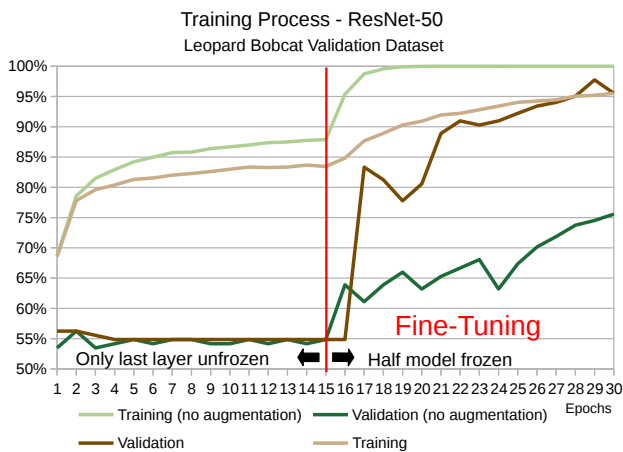
The red vertical bar marks the step between the different training phases for which different amounts of layers were frozen. The fine

tuning of the models significantly increased the performance of the models. The graph in Figure 8 shows an analysis on how the number of frozen layers during fine-tuning affects the performance. The number of layers was chosen to fit the residual blocks of the network architectures. When freezing more than half of the network's layer the performance stagnated. For the further analysis we froze the first half and retrained the second half of each model's layers in the fine-tuning process.

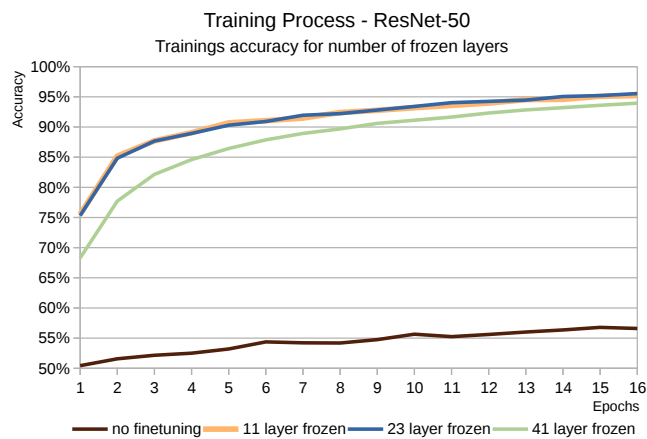
### 5.3 Performance in Case Study

Table 3 summarises the different models validated on the different validation datasets.

The fine-tuned models generally performed better than the basic models for which solely the rearmost layer of the network was retrained (Figure 3). On the lynx dataset with different light conditions, habitats and camera setups, the basic model predictions were poor with accuracies of lower than 50 %. The basic models performed better on the leopard and bobcat validation set, originating from the AP-10K dataset. Still, the predictions are too unreliable for real-world applications for camera trap analysis. The fine-tuning training phase significantly improved the performance of the models for unseen species (Figure 7). The accuracy doubled for most of the backbone models. The best performing model on the task of flank



**Fig. 7:** Performance of the model during the training process. The red line marks the switch from first to the second training phase for which less layers were frozen during transfer learning.



**Fig. 8:** During the training process different amounts of layers of the ResNet-50 were frozen. The performance stagnated for more than half of the layers being frozen.

prediction for the previously unseen lynx species in unknown habitat is the EfficientNetV2-S. It reached an accuracy of 88.70% compared to its competitors MobileNetV and ResNet-50 that reached 71.63% and 77.59% (Table 3).

All models performed better on the validation set with bobcats and leopards compared to the dataset of lynx. From the 135 images all models classified over 130 images correctly.

We assume that the characteristics of the dataset can explain the generally better performance for the bobcat and leopards compared to the lynx. The bobcat and leopard validation dataset was not included in the training phase, but the data originated from the AP-10K dataset [21] that was used during training. The images therefore resemble one another in terms of lighting, quality and environments, while the lynx dataset includes images captured in a different environment and included images during nighttime.

Images with snow in the background or images that include foreign objects like camera traps and collars tend to be classified incorrectly (Figure 10). If prey is present the flank prediction in some cases was done for the prey body. Another cause for incorrect predictions are animals in twisted positions (prediction highlighted with red box in Figure 6). The reason for the incorrect classification lies in the labeled data. The automatically extracted data labels do not cover all eventualities. In the above shown case the head and face are not visible and the front and hind paws are close together, for which the body keypoints from the classes for the front and the back of the animal cannot clearly be separated. We excluded such images in the label extraction process and therewith from the training dataset. Animals in such poses could not be properly learned by the models.

## 6 Conclusions

In this study we trained and compared flank prediction models for quadruped species using transfer learning on different backbone network architectures (Figure 3). The data labels were derived from datasets labeled on body keypoints for different mammalian species (Figure 2). To further improve the models, pose or flank labeled data during nighttime is required to make the models more robust in low light conditions. Data from a wider variety of habitats could make the model more robust towards new application cases.

The automatic label extraction was limited to the predefined conditions of the keypoints relative position to each other and their visibility in the image. Most cases were covered by one of the conditions, but some images of animals in twisted poses could not clearly be sorted into one of the classes of the visible flank. Examples include that one of the flanks was better visible than the other, but the position of the head and tail implied the opposite label or images where the animal is in unusual positions that do not meet the conditions made for an automatized extraction. The example presents



**Fig. 9:** Animals in twisted poses for which the conditions were not fulfilled or led to a wrong autolabelling of body keypoints.



**Fig. 10:** Misclassified images from the leopard and bobcat AP-10K dataset [21] and the lynx camera trap dataset.

a tiger vertically climbing a tree-like object. In this position the assumption of a quadruped gait is not met (Figure 9).

For future approaches the dataset could be extended with manually labeled data for twisted and unusual positions. Nevertheless, with animals being non rigid objects in versatile poses and habitats, predefined conditions or manually labeled data will not exhaustively cover all possibilities of animal poses for different species.

## References

- [1] Alexandra Swanson, Margaret Kosmala, Chris Lintott, et al. “Snapshot Serengeti, high-frequency annotated camera trap images of 40 mammalian species in an African savanna”. *Scientific data* 2 (2015), pp. 150026–150039. DOI: 10.1038/sdata.2015.26.
- [2] Stefan Schneider, Graham W. Taylor, Stefan Linquist, et al. “Past, present and future approaches using computer vision for animal re-identification from camera trap data”. *Methods in Ecology and Evolution* 10.4 (2019), pp. 461–470. DOI: 10.1111/2041-210X.13133.
- [3] Brett T. McClintock, Paul B. Conn, Robert S. Alonso, et al. “Integrated modeling of bilateral photo-identification data in mark-recapture analyses”. *Ecology* 94.7 (2013), pp. 1464–1471. DOI: 10.1890/12-1613.1.
- [4] Sonam Wangyel Wang and David W. Macdonald. “The use of camera traps for estimating tiger and leopard populations in the high altitude mountains of Bhutan”. *Biological Conservation* 142.3 (2009), pp. 606–613. DOI: 10.1016/j.biocon.2008.11.023.
- [5] Shuyuan Li, Jianguo Li, Hanlin Tang, et al. “ATRW: A Benchmark for Amur Tiger Re-identification in the Wild”. *MM '20: Proceedings of the 28th ACM International Conference on Multimedia* (2020), pp. 2590–2598. DOI: 10.1145/3394171.3413569.
- [6] Ben C. Augustine, J. Andrew Royle, Marcella J. Kelly, et al. “Spatial capture–recapture with partial identity: An application to camera traps”. *The Annals of Applied Statistics* 12.1 (2018). DOI: 10.1214/17-AOAS1091.
- [7] Daniel Joska, Liam Clark, Naoya Muramatsu, et al. “AcinoSet: A 3D Pose Estimation Dataset and Baseline Models for Cheetahs in the Wild”. *2021 IEEE International Conference (2021)*, pp. 13901–13908. DOI: 10.1109/ICRA48506.2021.9561338.
- [8] Allan F. O’Connell, James D. Nichols, and K. Ullas Karanth, eds. *Camera traps in animal ecology: Methods and analyses*. Springer, 2011. DOI: 10.1007/978-4-431-99495-4.
- [9] S. Palmero, J. Premier, S. Kramer-Schadt, et al. “Critical variables and their thresholds for the precise density estimation of wild felids with camera traps and spatial capture-recapture (SCR) methods”. *Mammal Review (in Press)* (2023).
- [10] Vanessa Suessle, Mimi Arandjelovic, Ammie K. Kalan, et al. “Automatic Individual Identification of Patterned Solitary Species Based on Unlabeled Video Data”. *World Society for Computer Graphics (WSCG) Proceedings (in Press)* (2023).
- [11] K. Weingarh, C. Heibl, F. Knauer, et al. “First estimation of Eurasian lynx (*Lynx lynx*) abundance and density using digital cameras and capture–recapture techniques in a German national park”. *Animal Biodiversity and Conservation* 35.2 (2012), pp. 197–207. DOI: 10.32800/abc.2012.35.0197.
- [12] *Zamba Cloud*. <https://www.zambacloud.com/> [Accessed: 15.06.2023]. 2023.
- [13] Devis Tuia, Benjamin Kellenberger, Sara Beery, et al. “Perspectives in machine learning for wildlife conservation”. *Nature communications* 13.1 (2022), pp. 792–806. DOI: 10.1038/s41467-022-27980-y.
- [14] Sara Beery, Dan Morris, and Siyu et al. Yang. “Efficient Pipeline for Automating Species ID in new Camera Trap Projects”. *Biodiversity Information Science and Standards* 3 (2019). DOI: 10.3897/biss.3.37222.
- [15] MMPose Contributors. *OpenMMLab Pose Estimation Toolbox and Benchmark*. <https://github.com/open-mmlab/mmpose>. 2020.
- [16] Shaokai Ye, Alexander Mathis, and Mackenzie W. Mathis. “Panoptic animal pose estimators are zero-shot performers” (2022). DOI: 10.48550/arXiv.2203.07436.
- [17] Jason R. Parham, Charles V. Stewart, Jonathan P. Crall, et al. “An Animal Detection Pipeline for Identification”. *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)* (2018), pp. 1075–1083. DOI: 10.1109/WACV.2018.00123.
- [18] Jinkun Cao, Hongyang Tang, Hao-Shu Fang, et al. “Cross-Domain Adaptation for Animal Pose Estimation”. In: *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*. 2019, pp. 9497–9506. DOI: 10.1109/ICCV.2019.00959.
- [19] Aditya Khosla, Nityananda Jayadevaprakash, Bangpeng Yao, et al. “Novel Dataset for Fine-Grained Image Categorization”. In: *First Workshop on Fine-Grained Visual Categorization, IEEE Conference on Computer Vision and Pattern Recognition*. 2011.
- [20] Alexander Mathis, Thomas Biasi, Steffen Schneider, et al. “Pre-training Boosts Out-of-Domain Robustness for Pose Estimation”. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*. 2021, pp. 1859–1868. DOI: 10.48550/arXiv.1909.11229.
- [21] Hang Yu, Yufei Xu, Jing Zhang, et al. *AP-10K: A Benchmark for Animal Pose Estimation in the Wild*. DOI: 10.48550/arXiv.2108.12617.
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, et al. “Deep Residual Learning for Image Recognition”. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016), pp. 770–778. DOI: 10.1109/CVPR.2016.90.
- [23] Mark Sandler, Andrew Howard, Menglong Zhu, et al. “MobileNetV2: Inverted Residuals and Linear Bottlenecks”. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2018), pp. 4510–4520. DOI: 10.1109/CVPR.2018.00474.
- [24] Mingxing Tan and Quoc Le. “EfficientNetV2: Smaller Models and Faster Training”. In: ed. by Marina Meila and Tong Zhang. Vol. 139. *Proceedings of Machine Learning Research*. PMLR, 2021, pp. 10096–10106.
- [25] *IUCN - Red List - Felidae*. <https://www.iucnredlist.org/search/list?taxonomies=101738&searchType=species> [Accessed: 22.06.2023]. 2023.
- [26] Stefano Palmero, Elisa Belotti, Luděk Bufka, et al. “Demography of a Eurasian lynx (*Lynx lynx*) population within a strictly protected area in Central Europe”. *Scientific reports* 11.1 (2021), pp. 19868–19878. DOI: 10.1038/s41598-021-99337-2.