# Towards Individual Grevy's Zebra Identification via Deep 3D Fitting and Metric Learning

Maria Stennett University of Bristol, UK gh18931@bristol.ac.uk Daniel I. Rubenstein Princeton University, USA dir@princeton.edu Tilo Burghardt University of Bristol, UK tilo@cs.bris.ac.uk

Abstract – This paper combines deep learning techniques for species detection, 3D model fitting, and metric learning in one pipeline to perform individual animal identification from photographs by exploiting unique coat patterns at below real-time. This is the first work to attempt this and, compared to traditional 2D bounding box or segmentation based CNN identification pipelines, the approach provides effective and explicit view-point normalisation and allows for a straight forward visualisation of the learned biometric population space. Note that due to the use of metric learning the pipeline is also readily applicable to open set and zero shot re-identification scenarios. We apply the proposed approach to individual Grevy's zebra (Equus grevyi) identification and show in a small study on the SMALST dataset that the use of 3D model fitting can indeed benefit performance. In particular, back-projected textures from 3D fitted models improve identification accuracy from 48.0% to 56.8% compared to 2D bounding box approaches for the dataset. Whilst the study is far too small accurately to estimate the full performance potential achievable in larger-scale real-world application settings and in comparisons against polished tools, our work lays the conceptual and practical foundations for a next step in animal biometrics towards deep metric learning driven, fully 3D-aware animal identification in open population settings. We publish network weights and relevant facilitating source code with this paper for full reproducibility and as inspiration for further research.

# I. INTRODUCTION AND BIOLOGICAL CONTEXT

Timely population data is a quintessential pillar required for most ecological studies and, critically, for effective conservation efforts [5]. In the past, accurate data on where and when individuals were seen was limited to only a few animals who were fitted with GPS tags or to those whose distinctive features easily stood out. Thus, inferences about population size, structure of dynamics associated with changes over time in population numbers, individual associations or habitats were limited because they were based on small samples. Biometric identification from photographs based on unique coat patterns avoids tagging and its computerisation addresses scaling issues. For plains zebras, for instance, symbolic stripe codes on manual filing cards were described by Klingel and later Petersen [6] in the mid-20th century. Computerised photographic identification started in the 1990s with first approaches by Hiby and Lovell [7], which explicitly normalised for viewpoint using semi-manual model fitting. Other early



Fig. 1: System Overview. We integrate four distinct deep learning systems into a fully automated pipeline for animal identification: (*blue*) animal detection and localisation via the MegaDetector [1]; (*green*) species recognition/confirmation via the Microsoft AI for Earth species classification API [2]; (*orange*) 3D zebra model fitting via the SMALST framework [3] and subsequent model back-projection of the visible texture portion of the coat pattern leading to a normalised surface texture; (*yellow*) Extracted textures are used to learn a latent population space via ResNet-driven metric learning [4]. Novel images projected into this space can be classified by k-nearest neighbour techniques, resulting in zebra identities.

ID methods either relied on user input or were limited to high quality [8], [9] or pre-processed imagery. The arrival of local descriptors [10] in combination with deep learning approaches [11] changed the field. Routine identification of individuals based on unique visual appearance of many species is now possible with good success for reasonable quality imagery [5], [12]. However, animal pose variations and deformations still cause challenges in individual re-identification despite the fact that metric learning [4] has opened new horizons for performance [13], open set applicability and population space visualisation [4]. Today, no framework exists that unifies modern approaches into an explicit 3D-aware deep metric learning setting. All existing automated animal ID techniques still rely on the implicit encoding of pose and viewpoint invariance within the classification framework or descriptor despite the availability of first explicit deep 3D fitting techniques for animals [3]. In response, we propose a first deep learning ID approach that integrates all desired aspects in a single pipeline (see Fig. 1). It is available at https: //github.com/Lm0079/grevys-zebra-individual-identification.

#### II. DATASETS

We use the 148 Grevy's zebra (Equus grevyi) images of the SMALST repository [3] for full system evaluation and combine it with 687 images from the WCS dataset [14] without individual labels to test species disambiguation. This SMALST+WCS dataset covers images of 300 Grevy's zebras, 232 plains zebras, 150 horses, and 153 donkeys. SMALST is one of the very few fully ID annotated and publicly available datasets for Grevy's zebra. It covers four images per individual across 37 distinct animals, two-thirds of which are male. Fig. 1 depicts sample images. Images were captured during the Great Grevy's Rally (GGR) in 2018 [15] over a two-day period at varying times and locations. Note that seeing the right side of a zebra provides no identifiable information on the left side and vice versa. Hence, all pictures taken of the zebras are of the right side. In addition, viewpoints are further labelled as right, front-right, and back-right. The train-test split for the images was generated using stratified sampling, so for each individual there are two images in the train and test set, respectively. To facilitate deep metric learning on this tiny train dataset, we applied significant offline augmentation, resulting in 50 training images per individual. Augmentations cover random colour jitters (saturation, contrast, and brightness), shifts, and rotation. In addition, we also use 15,000 synthetic images from SMALST [3] to train 3D fitting.



Fig. 2: MegaDetector Performance. (*left*) PR and ROC curve test performance for the MegaDetector on the SMALST dataset. Note the different, colour-coded units and abscissa crossings. (*right*) Example test images from SMALST of a true positive (TP) and a false negative (FN) detection. Whilst partial occlusions by vegetation are generally handled well, animal-animal occlusions often affect the detector's localisation ability negatively. This hints at a link between visual localisation in groups and dazzle coat pattern structure.



Fig. 3: Normalised Coat Pattern Extraction. The top row shows test image regions containing zebras with the MegaDetector [1] bounding box output (in red) and ground truth (in green). The middle row shows SMALST [3] fitted 3D zebra models superimposed on the 2D images. The bottom row depicts the back-projected and fully viewpoint and pose normalised texture maps of the zebra hindquarter and back used to identify the individual. Note that the middle two and left/right columns depict the same individual, respectively.

# III. METHODS AND COMPONENT PERFORMANCE *A. Animal Detection and 2D Localisation*

We used v4.1 of the state-of-the-art pre-trained MegaDetector [1] with its F-RCNN ensemble of Inception [16] and ResNet [17] streams to detect object instances of a generic 'animal' class first. We filter out any detections with a confidence score of less than 0.83 (determined via AUC optimisation), or if the detection is nested or nearly inside another (defined by IoU< 0.3). When evaluated on the SMALST dataset the MegaDetector showed a strong average precision (AP) of 0.96, which allows for region of interest (RoI) generation within the performance range expected for the MegaDetector [1] in different ecosystems. Fig. 2 (*left*) visualises the related performance curves and Fig. 2 (*right*) and Fig. 3 (*top row*) depict sample bounding box detections.

# B. Optional Species Disambiguation

For applicability to general imagery with visual distractor species, we use the pre-trained species classification model by Microsoft's AI for Earth team [2] to perform species disambiguation. This off-the-shelf, non-specific approach fuses Inception and ResNext outputs to confirm Grevy's zebra species presence against any other species. We filter out images where the top species prediction is not of a Grevy's zebra, irrespective of its species score and summarise results. Using this method on the full SMALST+WCS dataset accuracy was 81.3% at an F1 score of 0.88. This implies that even in high distractor cases species off-the-shelf disambiguation works fairly well for the species at hand. Note that this step is only required if random wildlife data (e.g. raw camera trap imagery) is fed into the system, whilst most studies apply ID systems to visual archives containing species of interest only.

#### C. Coat Texture Normalisation Approach

As bodies are deformable and viewpoints vary, changes in pose or camera angle alter the appearance of coat patterns. By aligning the 2D animal image with a 3D model and backprojecting the coat onto the fitted model surface the pose and viewpoint invariant texture can be estimated explicitly. Fig. 3 depicts test examples of this process. In particular, our pipeline feeds confirmed RoIs of Grevy's zebra to the SMALST [3] regression framework using SMAL horse templates [18], [19] to estimate the 3D shape, pose, and ultimately visible coat pattern surface texture.

#### D. Deep 3D Fitting Network

We follow the approach by Zuffi et al. [3] for training the texture fitting network. It consists of a ResNet backbone followed by a convolutional layer and 2 fully connected layers which produce a 1024-dimensional feature vector summarising appearance information of the input. This feature vector is fed to independent subsequent layers that predict shape, 3D pose, and texture maps. Shape prediction is implemented via fully connected layers that produce a 40-dimensional shape feature, used to predict vertex deformations applied to the SMAL 3D model. Pose prediction is performed by a linear layer that outputs a vector of relative joint angles. This vector sets the 3D pose of the SMAL model complemented by camera frame translation prediction via two linear layers that output coordinates (x,y,z). Finally, texture prediction uses encoderdecoder prediction of a uv-flow map from four stitched subimages [18]. We extract a compact subset from this map containing hindquarter and back areas to cover key unique portions of the coat (see Fig. 3 (bottom row)). Other characteristic areas such as scapular and neck regions could be added in further studies, but suffer from (self)occlusion and shadows more frequently.

#### E. Individual Identification via Metric Learning

We map cropped textures into learned, individually distinctive latent space following Andrew *et al.* [4], [20]. As a result, texture mappings of the same individual naturally cluster together. The space is built via triplets of inputs, which include an anchor pattern, a positive (pattern of the same individual as the anchor), and a negative (different individual



Fig. 4: Deep 3D Model Training. The graph depicts loss evolution during network training for 3D model fitting. Note that most losses components drop very quickly apart from the texture loss (top curve in light blue), which only gradually reduces with learning. Loss components on the right are explained in detail in [3]. An example of a fitted and textured 3D model is also shown.



TABLE I: 2D vs 3D Identification Performance. Top-1 accuracy results for three metric learning setups for 74 test images of 37 individuals of the SMALST dataset differentiating our 3D pipeline against two 2D baselines. Input examples are given at the bottom.

to the anchor). A ResNet model optimised via SoftMax and reciprocal triplet losses is used to learn the visual similarities and dissimilarities between the individuals forming an identity space (see Fig. 5 for a space visualisation). Unseen input textures can then be projected into this domain and a k-Nearest Neighbours approach will reveal closest individual identities.

#### **IV. ID EXPERIMENTS & RESULTS**

## A. 3D Fitting Training

We trained the 3D fitting network using 15,000 synthetic images from the SMALST training portion. Training proceeded for 210 epochs at a learning rate of 0.0001 and batch size 8 with a texture resolution of 256x256. Fig. 4 illustrates this training process quantitatively. We retained the network that performed best on the SMALST validation set in relation to the Percentage of Correct Keypoints (PCK)@0.1 and used it to extract normalised texture maps. Training ran for approx. 8 days on a BlueCrystal NVIDIA P100 node utilising one GPU.

#### B. Metric Space Training

Half of the images from SMALST are used for metric learning where each image is  $3 \times$  augmented increasing the training set size by factor 25. Further augmentations were not performed as it resulted in a plateau in performance. An online triplet mining method [21] known as 'batch hard' was used for sampling from this training set. We trained the network for 200 epochs employing a batch size of 8, a learning rate of 0.1, and a reciprocal triplet lambda [4], [20] of 0.0001. Training took approx. 4 hours utilising the hardware setup detailed above.

#### C. Basic Identification Experiments

We tested our pipeline against two simple baselines using SMALST data. First, we fed whole images without animal localisation into the metric learning component. Secondly, we used 2D texture patches of the zebra hindquarters and backs cropped from the 2D image by an F-RCNN localiser to assess in how far 2D data can compete with our 3D design. Results of our system versus these two baselines are shown in Table I and Fig. 5 visualises our identity space.



Fig. 5: Metric Identity Space of SMALST Population. We depict a 2D embedding of the 128D identity space generated by our pipeline covering the SMALST Grevy's zebra population containing 37 individuals. We use t-distributed Stochastic Neighbour Embedding (t-SNE) for dimensionality reduction. The colour code at the bottom maps to individual IDs left to right. Test data points carry a black outline, training data points are shown as simple colour discs. A sample pattern from the test set (bottom right) and two training patterns for a particular individual are also depicted. Note the very clean clustering of training information in this space. However, due to the tiny dataset the system overfits and various test data points are not reliably identified via the kNN approach. Future work will address these limitations.

### V. PROSPECTS AND CONCLUSION

This paper combined deep learning techniques for species detection, 3D model fitting, and metric learning to perform individual animal ID from photographs. This is the first work to attempt deep 3D ID. Whilst 2D computer vision algorithms changed the scale and scope individual-based ecological analyses, these techniques often require 2D images be filtered by pose, orientation, and clarity. Expanding image analysis to 3D provides significant advantages since fitting images to mannequin-like 3D poses reduces coat texture deformation, adds explicit pose information to the output, and can aid algorithmic accuracy as shown in this preliminary study. We applied our 3D approach to individual Grevy's zebra (Equus grevyi). We showed on the SMALST test dataset that the use of 3D model fitting can indeed benefit performance. However, the study is too small accurately to suppress overfitting and estimate the full performance potential achievable in largerscale real-world application. Large training data sets, higher resolution texture maps, and integration with existing non-deep frameworks will be needed to judge the pipeline w.r.t. stateof-the-art ID performance [4], [10], [13]. However, being able to construct and visualise population spaces based on posenormalised animal coat data and associate pose and ID information to this can benefit a range of behavioral and ecological studies. Further collaboration between the computer science and biological communities will be essential to enhance the data and algorithm base to improve on the quality of the biology that relies on re-sighting individuals.

Acknowledgements: We appreciate publications by GGR/SMALST, WCS, MegaDetector, AI4Earth, and Andrew/Lagunes. Thanks to T Berger-Wolf, C Stewart, and J Parham. This work was carried out using the computational facilities of the Advanced Computing Research Centre, University of Bristol - http://www.bris.ac.uk/acrc/.

#### REFERENCES

- [1] S. Beery, D. Morris, and S. Yang, "Efficient pipeline for camera trap image review," in *Data Mining and AI for Conservation at KDD*, 2019.
- Microsoft, "Species classification github repository," Accessed: 2022-06-24. [Online]. Available: https://github.com/microsoft/ SpeciesClassification
- [3] S. Zuffi, A. Kanazawa, T. Berger-Wolf, and M. J. Black, "Three-D Safari: Learning to Estimate Zebra Pose, Shape, and Texture from Images In the Wild," in *ICCV*, Oct. 2019.
- [4] W. Andrew, J. Gao, S. Mullan, N. Campbell, A. Dowsey, and T. Burghardt, "Visual identification of individual Holstein-Friesian cattle via deep metric learning," *Computers and Electronics in Agriculture*, vol. 185, 2021.
- [5] D. Tuia, B. Kellenberger, S. Beery, B. R. Costelloe, S. Zuffi, B. Risse, A. Mathis, M. W. Mathis, F. van Langevelde, T. Burghardt *et al.*, "Perspectives in machine learning for wildlife conservation," *Nature communications*, vol. 13, no. 1, pp. 1–15, 2022.
- [6] J. Petersen, "An identification system for zebra (Equus burchelli, gray)," African Journal of Ecology, no. 10, pp. 69–63, 1972.
- [7] L. Hiby and P. Lovell, "Computer aided matching of natural markings: a prototype system for grey seals," *Report of the International Whaling Commission*, no. 12, pp. 57–61, 1990.
- [8] 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 ACM International Conference on Multimedia Retrieval, 2011, pp. 1–8.
- [9] R. B. Sherley, T. Burghardt, P. J. Barham, N. Campbell, and I. C. Cuthill, "Spotting the difference: towards fully-automated population monitoring of african penguins spheniscus demersus," *Endangered Species Research*, vol. 11, no. 2, pp. 101–111, 2010.
- [10] J. Crall, C. Stewart, T. Berger-Wolf, D. Rubenstein, and S. Sundaresan, "Hotspotter - patterned species instance recognition," in 2013 IEEE Workshop on Applications of Computer Vision, 2013, pp. 230–237.
- [11] J. Parham, C. Stewart, J. Crall, D. Rubenstein, J. Holmberg, and T. Berger-Wolf, "An animal detection pipeline for identification," in WACV, 2018, pp. 1075–1083.
- [12] T. Y. Berger-Wolf, D. I. Rubenstein, C. V. Stewart, J. A. Holmberg, J. Parham, S. Menon, J. Crall, J. V. Oast, E. Kiciman, and L. Joppa, "Wildbook: Crowdsourcing, computer vision, and data science for conservation," in *Bloomberg Data for Good Exchange Conference*, 2019.
- [13] S. Schneider, G. W. Taylor, and S. C. Kremer, "Similarity learning networks for animal individual re-identification-beyond the capabilities of a human observer," in *Proceedings of the IEEE/CVF Winter Conf. on Applications of Computer Vision Workshops*, 2020, pp. 44–52.
- [14] W. C. Society, "Wildlife conservation society (WCS) camera traps," Accessed: 2022-06-30. [Online]. Available: https://lila.science/datasets/ wcscameratraps
- [15] D. Rubenstein, J. Parham, C. Stewart, T. Y. Berger-Wolf, J. Holmberg, J. Crall, B. L. Mackey, S. Funnel, K. Cockerill, Z. Davidson, L. Mate, C. Nzomo, R. Warungu, D. Martins, V. Ontita, J. Omulupi, J. Weston, G. Anyona, G. Chege, D. Kimiti, K. Tombak, A. Gersick, and N. Rubenstein, "The State of Kenya's Grevy's Zebras and Reticulated Giraffes: Results of the Great Grevy's Rally 2018," Accessed: 2022-06-30. [Online]. Available: https://wildlifedirect.org/wp-content/ uploads/2018/06/CS-report-GGR-2018v-4.pdf
- [16] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in *CVPR*, 2015, pp. 1–9.
- [17] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in 2016 IEEE Conf. on Computer Vision and Pattern Recognition, 2016, pp. 770–778.
- [18] S. Zuffi, A. Kanazawa, D. Jacobs, and M. J. Black, "3D menagerie: Modeling the 3D shape and pose of animals," in *IEEE Conf. on Computer Vision and Pattern Recognition*, Jul. 2017.
- [19] S. Zuffi, A. Kanazawa, and M. J. Black, "Lions and tigers and bears: Capturing non-rigid, 3D, articulated shape from images," in CVPR, 2018.
- [20] M. Lagunes-Fortiz, D. Damen, and W. Mayol-Cuevas, "Learning discriminative embeddings for object recognition on-the-fly," in 2019 International Conf. on Robotics and Automation. IEEE, 2019, pp. 2932–2938.
- [21] F. Schroff, D. Kalenichenko, and J. Philbin, "FaceNet: A unified embedding for face recognition and clustering," in 2015 IEEE Conf. on Computer Vision and Pattern Recognition (CVPR). IEEE, jun 2015.