

Animal3D: A Comprehensive Dataset of 3D Animal Pose and Shape

Jiacong Xu¹ Yi Zhang¹ Jiawei Peng¹ Wufei Ma¹ Artur Jesslen¹¹ Pengliang Ji³
 Qixin Hu⁴ Jiehua Zhang⁵ Qihao Liu¹ Jiahao Wang¹ Wei Ji⁶ Chen Wang⁷
 Xiaoding Yuan¹ Prakhar Kaushik¹ Guofeng Zhang⁸ Jie Liu⁹ Yushan Xie²
 Yawen Cui⁵ Alan Yuille¹ Adam Kortylewski^{10,11}

¹Johns Hopkins University

²East China Normal University

³Beihang University

⁴HUST ⁵University of Oulu

⁶University of Alberta

⁷Tsinghua University

⁸UCLA

⁹City University of Hong Kong

¹⁰Max Planck Institute for Informatics

¹¹University of Freiburg

Abstract

Accurately estimating the 3D pose and shape is an essential step towards understanding animal behavior, and can potentially benefit many downstream applications, such as wildlife conservation. However, research in this area is held back by the lack of a comprehensive and diverse dataset with high-quality 3D pose and shape annotations. In this paper, we propose Animal3D, the first comprehensive dataset for mammal animal 3D pose and shape estimation. Animal3D consists of 3379 images collected from 40 mammal species, high-quality annotations of 26 keypoints, and importantly the pose and shape parameters of the SMAL [50] model. All annotations were labeled and checked manually in a multi-stage process to ensure highest quality results. Based on the Animal3D dataset, we benchmark representative shape and pose estimation models at: (1) supervised learning from only the Animal3D data, (2) synthetic to real transfer from synthetically generated images, and (3) fine-tuning human pose and shape estimation models. Our experimental results demonstrate that predicting the 3D shape and pose of animals across species remains a very challenging task, despite significant advances in human pose estimation. Our results further demonstrate that synthetic pre-training is a viable strategy to boost the model performance. Overall, Animal3D opens new directions for facilitating future research in animal 3D pose and shape estimation, and is publicly available.

1. Introduction

Accurately estimating the 3D pose and shape of animals is a crucial step toward understanding their behavior and has a wide range of applications in fields such as wildlife conservation, animal ecology, and biomechanics. 3D ani-

mal pose and shape estimation involves the reconstruction of the 3D structure of an animal from a single 2D image, which is a challenging task due to the complex shapes and poses of animals in the wild. Previous works in this area have primarily focused on specific animals, such as humans [16, 17] or dogs [37, 45], which limits the generalization ability of the models to other animals. Therefore, there is a need for a diverse dataset of animals to allow for more generalizable and robust models to be developed.

In this paper, we propose Animal3D, the first benchmark for mammal animal 3D pose and shape estimation. Animal3D is a comprehensive dataset consisting of 3379 high-quality images collected from 40 mammal species. The images were carefully selected from existing datasets, in particular PartImageNet [11] and COCO [24], to ensure that they represent a diverse range of animals, including primates, ungulates, carnivores, and rodents (Figure 1). The diversity of animal species included in the dataset ensures that the models are not limited to specific animals and can be applied to a wide range of species. Each image was annotated with 26 keypoints, which were carefully labeled and checked in a multi-stage process to ensure high-quality annotations that can be used for further research. Based on the keypoint annotation and the available segmentation masks in PartImageNet and COCO, we annotate the 3D shape and pose by fitting the SMAL [50] model to the data. SMAL is a widely used model for 3D animal pose and shape estimation and similarly structured as the SMPL [26] model, hence supporting wide applicability of our annotations.

Using the Animal3D dataset, we benchmark representative shape and pose estimation models at three levels: (1) supervised learning for animal pose estimation, (2) synthetic to real transfer from synthetically generated images, and (3) fine-tuning human pose and shape estimation models. Based on our experimental results we provide an anal-

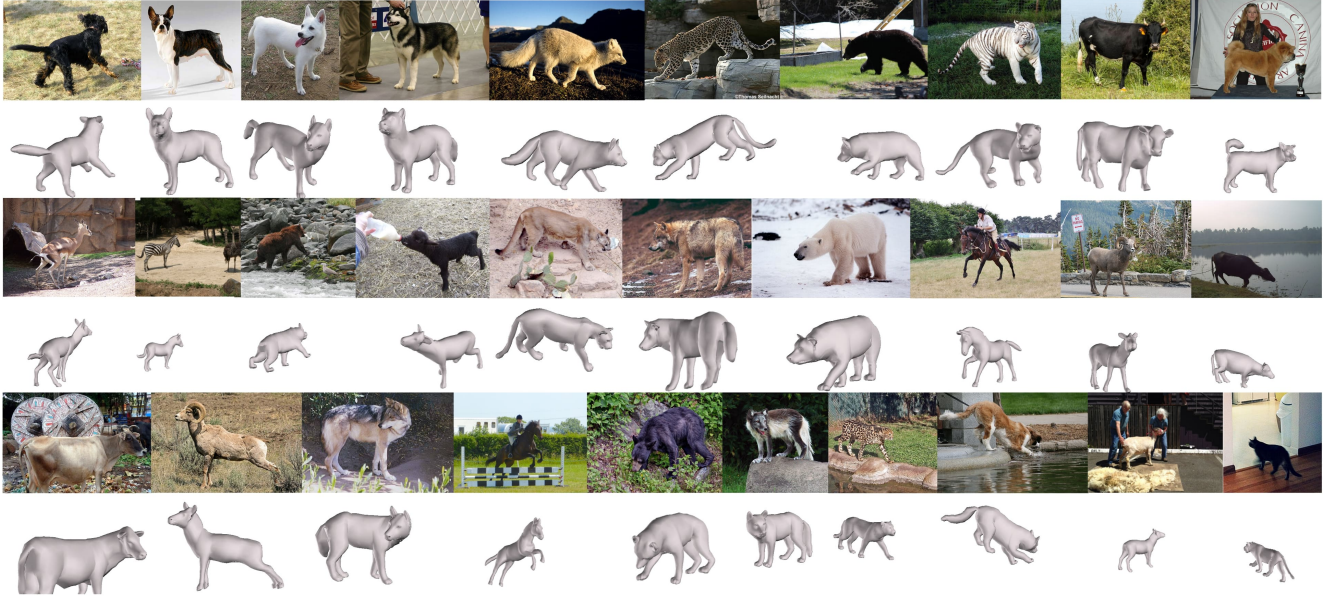


Figure 1: Samples from the proposed Animal3D dataset. Our dataset contains a diverse range of animal species with high-quality annotations of shape and pose parameters using the popular SMAL [50] model.

ysis of the strengths and limitations of each method, which demonstrate the versatility of our benchmark, as well as its challenging nature, since none of the representative approaches achieves a similarly good performance as on the specialized benchmarks they were designed for.

Animal3D is a significant step towards facilitating future research in animal 3D pose and shape estimation. The dataset will allow researchers to advance the understanding of animal behavior and ecology through 3D pose and shape estimation. Additionally, the dataset has the potential to benefit many downstream applications. The models developed using the Animal3D dataset can be applied to a wide range of animals, potentially leading to new discoveries and insights into animal behavior and ecology. Access the data via: <https://xujiacong.github.io/Animal3D>

In summary, our main contributions are:

- We present Animal3D, the first benchmark for mammal animal 3D pose and shape estimation, with a diverse set of 40 mammal species, and high-quality annotations of 2D keypoints as well as 3D shape and pose parameters of the SMAL [50] model.
- We set up a set of baselines on Animal3D in various settings using state-of-the-art methods which demonstrates the versatility of the dataset.
- Our experimental results and in-depth analysis of the strengths and limitations of representative methods demonstrate the challenging nature of our benchmark.

2. Related Work

Existing animal datasets and the development and challenges of pose estimation methods for both human and animals will be discussed in this section.

2.1. Animal Pose Estimation

Datasets. Existing animal datasets provide 2D annotations, such as keypoints, bounding box, or segmentation, on single or multiple species. AP-10K [46] is the largest public dataset that contains over 10k images for 54 animal species and 17 keypoints are defined for each animal. Horse-30 [27] comprises 8k frames for 30 different horses and 22 keypoints are annotated for each image. AcinoSet [15] records 119k multi-view frames for Cheetahs and annotates 20 keypoints on 7.6k frames. StanfordExtra [2] extracts 8.1k images from Stanford dogs and annotates 20 keypoints and segmentation on many dog breeds. Animal Pose Dataset [6] extends the Poselets dataset by annotating more images from Animal-10 and recently serves as the standard benchmark for 2D animal pose estimation due to its large scale and rich species. Unlike human pose datasets, there is no widely accepted annotation standard for animals since the biological diversity between different animal species is much more significant than among humans. For example, the keypoints of knee in Animal Pose Dataset are defined as middle points on the leg in StanfordExtra, and thereby their positions are slightly different. Nevertheless, 2D keypoint detection is not enough to derive the full geometry information of the objects and there is no known dataset that

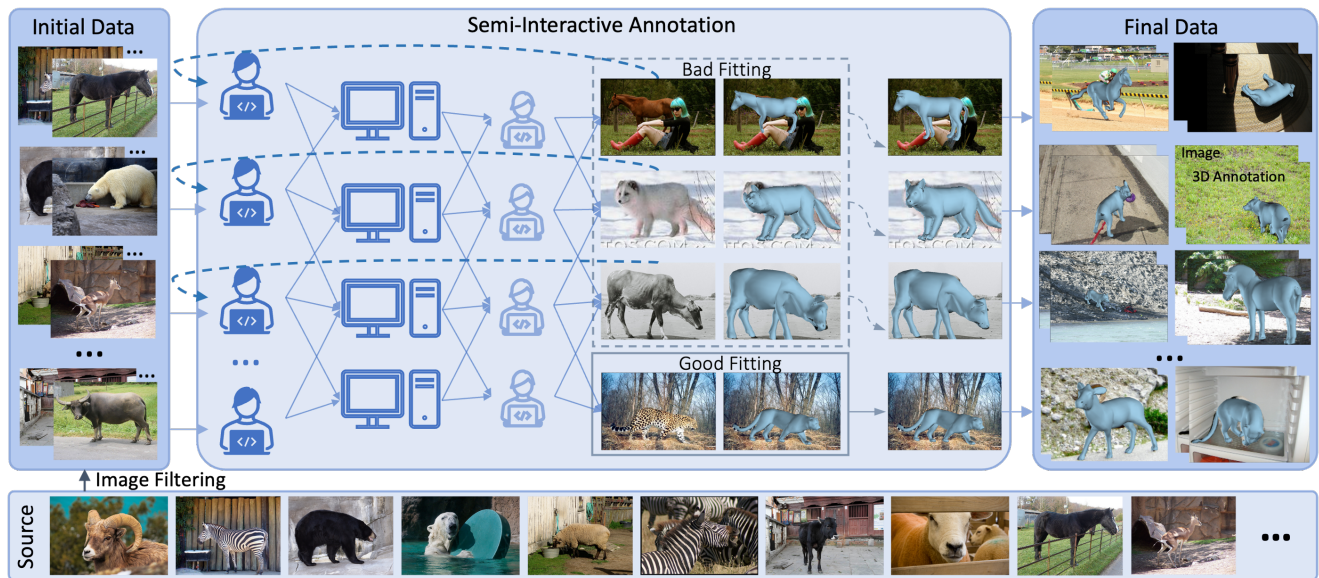


Figure 2: Data annotation pipeline for Animal3D. The process consists of three stages: Image Filtering, Semi-Interactive Annotation, and Data Integration. The data is sourced and filtered to obtain an initial set of images. During the Semi-Interactive Annotation, annotators submitted their annotation to the server to fit the SMAL model and render the results on the images. Then a set of inspectors examined the fitting results and send the bad-fitting images back to the annotator for revision. This process is repeated multiple times. Images that constantly lead to bad-fitting results are removed.

contains 3D annotation for animals.

Methods. Large amount of data is required for better performance of deep models, but the collection and annotation processes of animal images are much more complex compared with humans. To deal with the data scarcity of animals, researchers came up with many efficient and effective ways. For instance, Cao et al. [6] feed the models with animal and human data together and employ domain adaptation to align the feature projection space. Mu et al. [30] construct a synthetic animal dataset with different textures and poses and utilize the unlabeled images with generated high-score pseudo labels to train the model. To further reduce the domain gap between synthetic and real data, Li and Lee [23] feed the multi-scale information to the domain classifier with gradient reverse layer [8]. These models focus on the 2D keypoint detection task of animals, where benchmarks are available for a fair comparison.

Nevertheless, the research on 3D pose estimation of animals is proceeding slowly due to the vacancy of available dataset. By mapping the pixels to vertices on template model (CSM), Kulkarni et al. [21] introduce an learning-based approach to optimize the articulation of the template. LASSIE [45] is the first work to recover the shape of articulated object without any template or prior models. The Skinned Multi-Animal Linear (SMAL) model proposed by Zuffi et al. [51] built a parametric way to represent the animal shape and pose based on strong prior and its modeling

performance are much better than non-parametric methods.

Biggs et al. [4] replace the manual labeling process for keypoints and silhouettes in SMAL by the prediction of pre-trained deep CNNs. By joint optimization on multiple images for the same animal, SMALR [49] recovers more shape and texture details. SMALST [48] directly regresses the shape and pose parameters of SMAL and textures for Zebras, and utilizes the difference between rendered and original images to optimize the neural features. WLDO [2] and BARC [37] are built based on an adapted version of SMAL for dogs and achieve satisfactory 3D recovery performance on various dogs. Since there is no dataset with 3D annotation, aforementioned works can only be evaluated using 2D measures, like 2D keypoints or masks.

2.2. Human Pose Estimation

In contrast to animal pose estimation, human pose has been studied extensively in the computer vision literature. Regression-based methods [10, 16, 32, 35, 39, 40] directly estimate 3D human pose from RGB image using a deep network. Different 3D human pose representations are adopted such as 3D joint locations [28, 36], 3D heatmaps [34, 38, 47] and parameters of a parametric human body [16, 35, 18]. To model pose ambiguities, e.g. for truncated human images, [3, 20] predict multiple possible poses that have correct 2D projections. Optimization-based methods [33, 19, 5, 42, 44] involve parametric human mod-

els like SMPL [33, 1, 26], and produce both the 3D human pose and human shape. The representative method is SMPLify [5], which fits the SMPL model to 2D keypoint detections with strong priors. Exploiting more information into the fitting procedure has been investigated, including silhouettes [22], multi-view [12], more expressive shape models [14]. [42] propose to fit 3D part affinity maps to overcome 2D-3D ambiguity.

While there has been a lot of progress in human pose estimation, most methods require large-scale annotated training data and therefore it remains unclear if these approaches can generalize to other animal species.

3. Animal3D Dataset

In this section, we introduce the Animal3D benchmark and discuss the data collection and annotation processes.

Existing annotation methods for 3D human pose estimation datasets [13, 29, 41], which utilize wearable devices, laser body scanner, multi-camera studio to capture the accurate motion and shape of the humans, cannot be generalized to animals since animals are not as controllable as humans and some are even dangerous. To still enable 3D pose and shape annotation of animals, we follow the interactive keypoint annotation idea used in PASCAL3D+ [43] for 3D pose annotation of rigid objects. In particular, we implement a web-based annotation tool that enables fast and accurate keypoint annotations. We extend it to articulated animals by using the SMAL [51] to fit the 3D animal model to annotated keypoints and segmentation masks. Overall, we manually collect and annotate images first, and subsequently conduct three rounds of quality checking and revision process, as illustrated in Figure 2. Compared to other animal datasets, Animal3D is the first dataset that provides 3D annotation of animals (Table 1).

	Animal3D (Ours)	Animal Pose[6]	Stanford Extra[2]	AP-10K [46]
Segmentation	✓	✗	✓	✗
3D Anno.	✓	✗	✗	✗
#Species	40	5	Dogs	54
#Keypoints	26	20	20	17
#Images	3.4K	4K	8.1K	10K

Table 1: Comparison of Animal3D with other animal datasets. Animal3D contains class labels of 40 species, 26 keypoints, and 3D pose and shape parameters from the SMAL model. Totally, there are 5.1k images are carefully annotated in Animal3D, but only 3.4k images are selected after 3-round inspection. The unselected images and annotations will also be published together with Animal3D.

3.1. Data Collection

Source. Our aim is to obtain shape and pose parameters by fitting SMAL to images using keypoints and silhouette (foreground segmentation) annotations as described in Section 3.3. To simplify the annotation process, we source the animal images and segmentation masks from existing datasets. After investigation of the existing segmentation datasets, we choose PartImageNet [11] and COCO [25] as our source datasets since they provide accurate segmentation masks of a diverse set of animals.

Filtering. Due to the limitation of the PCA shape space of SMAL, some of the animals cannot be represented properly by the SMAL model, such as the elephants and giraffes. Therefore, we remove images belonging to these categories. Additionally, there are a large number of images in which the animals are highly occluded or truncated. This sometimes makes it even challenging for humans to guess the invisible animals’ pose or parts correctly. Therefore, we also removed these images from the data. Finally, we selected a total of 5.1k images of 40 mammal species. Details on the exact animal classes and and image statistics can be found in the supplementary material.

Animal class labels. Unlike PartImageNet where all the images are grouped into ImageNet [7] categories, COCO does not provide fine-grained class labelling. To preserve a detailed category-level annotation, we manually classified the images from COCO into ImageNet categories.

3.2. Data Annotation

Since the pose and shape deformation of the SMAL model is highly dependent on the 2D keypoint annotation, we annotate 26 keypoints per animal based on the original keypoint definition of SMAL model (Figure 3). For the keypoint annotation, an interactive is important to guarantee the consistency of the keypoint locations across different images for different animal species, and to make the SMAL fitting results as precise as possible. However, the fitting and rendering of the SMAL model cannot be implemented in real-time therefore a fully interactive annotation was not possible. Instead, we designed an semi-interactive pipeline to make the annotation process as interactive as possible, as described in the following.

Annotators. Each keypoint annotator were assigned approximately 300 images belonging to 3-5 similar animal species. They were suggested to start with a few simple cases, where the entire body of the animal is able to be seen clearly, to become familiar with the corresponding animals. For invisible keypoints caused by occlusion, the annotators were asked to guess the positions of the keypoints based on their annotation experience and mark them as invisible.

Inspectors. The annotation inspectors examine the SMAL fitting results based on the initial annotation to ensure a high annotation quality. They also conduct extra ac-



Figure 3: Visualization of the 26 keypoints that are annotated in the Animal3D model. Other popular 2D datasets only annotate the visible keypoints, while we ask the annotators to guess the location of occluded or truncated body parts based on their annotation experience, which significantly improves the fitting performance of SMAL model.

tions to further improve the fitting results. For example, they compare the results with and w/o invisible keypoints, or change hyper-parameters of the SMAL fitting process.

Annotation Pipeline. During the annotation process, the annotators send their annotations to the servers to fit the SMAL model and render the results. Subsequently, the inspectors assess the annotation quality, filtering out the good-fitting examples and from the bad-fittings. If the bad fitting is caused by the keypoint annotation, the inspectors send the images back to the annotators for revision, and provide feedback on where the annotation can be improved. If the bad fitting is caused by a broken segmentation mask from slightly occluded or truncated objects, the inspectors gradually decrease the weights of silhouette error in the objective function to reduce the effect of mask, which in these cases typically improves the fitting results significantly. This annotate-then-examine process will proceed for three rounds. After the final round, the remaining bad-fitting images were discarded.

3.3. Fitting Animals to Images

The SMAL model $M(\beta, \alpha, t)$ is a function of shape β , pose α , and translation t . We fit the model to images by optimizing the model parameters guided by a combination of 2D keypoints and 2D silhouettes, as proposed in [50], with minor modifications. In the following, we provide a concise description of the fitting process, more details can be found in the original work. We denote $P(\mathbf{v}_i)$ as the perspective projection of the the i 'th mesh vertex into the image plane. Moreover, $P(M) = S$ is the projected mesh silhouette. To fit SMAL to an image, we optimize the model parameters $\Pi = \{\beta, \alpha, t\}$ over a joint loss function that is composed of the reprojection error of the keypoint, the silhouette repro-

jection error, and a shape prior

$$\mathcal{L}_{total}(\Pi, M) = \mathcal{L}_{kp}(\Pi, M) + \mathcal{L}_{silh}(\Pi, M) + \mathcal{L}_{shape}(\beta). \quad (1)$$

The losses are weighted to be approx. in the same range.

Keypoint loss. Each keypoint on the SMAL model corresponds to a subset of the mesh vertices. We denote this set of keypoint vertices as $\mathbf{v}_j^i \in V_i$, and optimize their projection to match the corresponding annotated keypoint t_i :

$$\mathcal{L}_{kp}(\Pi, M) = \sum_{i=1}^{26} \rho\left(\left\|\sum_{j=1}^{N^i} P(\mathbf{v}_j^i)/N^i - t_i\right\|_2\right), \quad (2)$$

where ρ is the Geman-McClure robust error function [9] to reduce the negative effects of difficult to fit annotations.

Silhouette loss. The silhouette is optimized using a bi-directional distance:

$$\mathcal{L}_{silh}(\Pi, M) = \sum_{x \in S} \mathcal{D}_{\bar{S}}(x) + \sum_{x \in \bar{S}} \rho\left(\min_{\hat{x} \in S} \|x - \hat{x}\|_2\right), \quad (3)$$

where \bar{S} is the ground truth silhouette and $\mathcal{D}_{\bar{S}}$ is its distance transform. The weight for the first term are manually adjusted to reduce the effect of occlusion and truncation.

Shape prior. We regularize the shape parameters β using a shape loss $\mathcal{L}_{shape}(\beta)$ using the PCA prior distribution. In particular, the loss is defined to be the squared Mahalanobis distance defined using the PCA eigenvalues.

3.4. Data Summary

Building on previous datasets and our additional annotation, the Animal3D dataset presents a comprehensive set of annotations for each image, including detailed Imagenet class labels, segmentation masks, 26 keypoints, and SMAL parameters for shape, pose and translation. Therefore, Animal3D can serve as a benchmark for a number of tasks as well as multi-tasking, 3D reconstruction or synthetic to real domain adaptation. We believe that our dataset will enable significant advances in all of these research areas, due to its high-quality annotation and scale.

4. Experiments

In this section, we benchmark representative shape and pose estimation models on Animal3D in three settings: (1) supervised learning (Section 4.1), (2) synthetic to real transfer from synthetically generated images (Section 4.2), and (3) fine-tuning pre-trained human pose and shape estimation models (Section 4.3).

Baselines. As we present the first comprehensive dataset of 3D shape and pose annotations for animals, there are no baselines that were explicitly designed for animal pose estimation in such a diverse setting. Nevertheless, strong representative baselines exist for human pose estimation and

Method	Supervised			Synthetic to Real			Human pre-trained		
	PA-MPJPE↓	S-MPJPE↓	PCK↑	PA-MPJPE↓	S-MPJPE↓	PCK↑	PA-MPJPE↓	S-MPJPE↓	PCK↑
HMR [16]	140.7	496.2	59.3	124.8	497.7	63.1	132.2	488.0	60.6
PARE [17]	134.8	443.9	79.1	127.2	392.3	83.7	130.7	374.9	85.6
WLDO [2]	128.8	502.1	60.1	123.9	484.0	65.1	-	-	-

Table 2: 3D pose and shape estimation results on the Animal3D dataset. We evaluate three representative baseline models, HMR, PARE and WLDO, in three settings: (1) Supervised on Animal3D data only, (2) Pre-training on synthetic data and fine-tuning on Animal3D, and (3) Pre-training on Human Pose Estimation datasets and fine-tuning on Animal3D. While pre-training improves results for all models, the final results are lower compared to object specific benchmarks for humans and dogs, hence indicating the difficulty of estimating 3D animal pose across species.

for 3D pose estimation of specific animal classes, which we adapt to the Animal3D dataset. We chose HMR [16] and PARE [17] as competitive and robust baselines for human pose estimation. Moreover, we selected WLDO [2] as a baseline that was specifically designed for animals, although only for dogs.

Evaluation metrics. We report scale-aligned mean per joint position error (S-MPJPE) and Procrustes-aligned mean per joint position error (PA-MPJPE) in mm as the main evaluation metrics, where the latter is the former plus rotational alignment. We do not use the popular per joint position error (MPJPE) in 3D human pose estimation since the scale of animals can vary a lot. We also report the 2D Percentage of Correct Keypoints (PCK) with threshold defined by half of the head-to-tail length to measure how well the prediction aligns with the 2D image.

Model & Data Preparation. We split the Animal3D dataset into 3059 training and 320 test images, by randomly sampling from the full dataset.

For training HMR [16], we remove its discrimination loss and keep only the 2D and 3D supervision, since there is no fake and real pose parameter available. PARE [17] requires the part grouping of the model vertices to render the ground-truth of 2D part segmentation. To obtain these labels, we manually segment all the vertices on SMAL into the 7 object parts defined in PartImageNet [11]: Head, Torso, Tail, and the four legs. To train WLDO [2], we remove the EM process that was designed to deal with different dog families, and we supply WLDO with direct supervision of shape and pose parameters for fair comparison with the human models.

Training setup. To be consistent with existing training pipelines for human pose estimation models, the animals are cropped from the image based on their bounding box and resized to 224×224 . Random rotation and flipping are implemented for data augmentation. The batch size for all the experiments are set to be 128 and we trained the models on 2 RTX3090 GPUs using synchronized batch normalization. For experiments with synthetic data, we pre-train the

models for 100 epochs. For training on real data, all the models are trained for 1000 epochs (around 24k iterations).

4.1. Supervised Animal Shape and Pose Regression

The left part of Table 2 shows the results of training HMR, PARE and WLDO in a supervised manner only from images of the Animal3D dataset. The ranking among the methods is as expected. HMR, which is an older method, performs worse compared to recently developed PARE model. Although the performance gap in PA-MPJPE is smaller compared to the respective results on human pose estimation datasets. WLDO performs best in terms of PA-MPJPE, hence suggesting that it predicts the articulation of the animals most accurately. However it does not perform particularly well at predicting the rigid 3D body pose, hence it achieves the worst results in terms of S-MPJPE. Notably, we observe that the prediction accuracy of all baseline models is significantly lower compared to their performance on the original domains that they were initially designed for. Hence, pointing out that the **3D animal pose estimation problem remains an important open research problem**. We believe that this performance gap is caused mainly by two main factors. The lack of large-scale annotated datasets, and the architectural design of the baseline methods for a particular object class, i.e. humans and dogs. In the following we aim to address the data problem using synthetic data and pre-training on large scale human data.

4.2. From Synthetic to Real

For human pose estimation, a larger amount of high-quality data usually will lead to better regression performance. Nevertheless, the annotation process for Animal3D is too complex to make it a larger dataset, so we are seeking more convenient way to generate more usable data. Inspired by [31], which utilize rendered images from CAD model of animals to boost the model performance on 2D tasks, we synthesize 45k images using SMALR [49] and select 40k for training and 5k for testing models, respectively, before fine-tuning them on Animal3D.




Figure 4: Example images from our synthetic dataset that is used for pre-training the animal pose estimation baselines. We simulate all species from the Animal3D dataset using the SMALR model in varying poses, shapes, and background images.

For each class in our dataset, several unoccluded and non-truncated images are selected to fit the SMALR model for the basic shape and texture. Then, we calculate the mean and covariance of the shape and pose parameters in the training set and sample the parameters from multi-variate Gaussian to mimic the realistic shape and pose for corresponding animal category. The background images are randomly selected from ImageNet [7] and consist of indoor and outdoor scenarios. Here are some examples for the synthetic data in Figure 4. Note that, the generation process of the synthetic data is based on the prior information obtained from Animal3D dataset, so its searching space can be enlarged by moderate the covariance. The center part of Table 2 shows that all methods benefit from synthetic pre-training. PARE benefits the most on average across metrics, outperforms HMR and WLDO significantly in terms of S-MPJPE and PCK, and thereby shows the potential of synthetic pre-training for animal pose estimation.

4.3. From Human to Animal

A common approach to training deep networks in a data efficient manner, is to initialize with models that are pre-trained on large datasets in related tasks. We study the effect of using pre-trained human pose estimation models as initialization to train an animal pose estimator. Both HMR and PARE have been trained on large-scale human data including Human3.6M [13], MPI-INF-3DHP [28], and COCO [25] datasets, and we use the publicly available models to fine-tune them on Animal3D.

The right part of Table 2 shows that both models outperform their non-pre-trained counterparts. Interestingly, the performance gap between HMR and PARE model in terms of S-MPJPE and PCK increases due to human pre-training. However, compared to the pre-training on synthetic data, which is much more easy to achieve, pre-training on real human data does not show a benefit.

4.4. Qualitative Results

Figure 5 illustrates qualitative regression results of several models that we have tested in Table 2. We observe that human-pretrained models always fail to recover the shape information of the animals and generate some unrealistic shapes. On the contrary, the models pretrained by synthetic data regress the shape parameters better. We argue that the model is able to learn strong shape prior from synthetic data, while it will focus more on the pose information for human data since the shape diversity between humans are much smaller than among animals. Also, the domain gap between human and animals and different feature projection space of SMPL [26] and SMAL will hinder the generalization of the models. Even so, in most cases the shape parameters are approximately correctly estimated, i.e. often the correct animal species is predicted. However, the alignment of the legs and the gaze direction are often incorrect. Overall, PARE demonstrates also qualitatively that its predictions have the best quality, mostly because they align better to the image, which can also be observed from its high PCK.

4.5. Discussion

Based on our results, we observe that PARE is a promising model for 3D animal pose estimation, as it achieves the best performance among the representative baselines. Its high performance is very likely caused by its advanced architecture that uses additional supervision in terms of part segmentations. In terms of scaling deep learning to animal 3D pose estimation, our results show that pre-training on large-scale synthetic data is a promising direction forward. Nevertheless, we observe that none of the baselines achieves a satisfying performance compared to the results obtained on the specific domains that the baselines were originally designed for. Hence, our experimental results demonstrate that predicting the 3D shape and pose of animals across species remains a very challenging task, despite significant advances in human pose estimation and animal pose estimation for specific species.




Figure 5: Visualization of regression performance of human and animal pose estimation models. The columns from left to right refer to the input image, the groundtruth from Animal3D, regression results for HMR, HMR pretrained by synthetic and human data, PARE and WLDO pretrained by synthetic data, respectively.

5. Conclusion

Animal3D is unique and diverse in that it includes 3D annotations for a large number of animal species, making it the first benchmark for mammal animal 3D pose and shape estimation. The comprehensive nature of Animal3D, in terms of diversity of animals and annotations for multiple vision tasks (keypoints, 3D SMAL parameters, segmentation) provides a foundation for the development of more robust and generalizable models for animal 3D pose and shape estimation. Animal3D and the implementation of the baselines will be publicly available, and we encourage researchers to use it for further research or applications.

The results of our experiments demonstrate that Animal3D is a valuable resource for improving animal pose and shape estimation models. We show that our dataset can be used to benchmark supervised learning for animal

pose estimation, synthetic to real transfer, and fine-tuning human pose and shape estimation models. We observe that existing methods for human pose estimation, achieve competitive results at animal pose estimation, when pretrained on synthetic data. However, the prediction performance is significantly lower compared to their accuracy on human-specific benchmarks. These results highlight that the 3D animal pose estimation task remains an important open research problem. These experiments will provide a strong foundation for future research in this area, which will benefit both scientific understanding and conservation efforts.

Acknowledgements. Adam Kortylewski acknowledges support via his Emmy Noether Research Group funded by the German Science Foundation (DFG) under Grant No. 468670075. Alan Yuille acknowledges support from Army Research Laboratory award W911NF2320008 and Office of Naval Research N00014-21-1-2812.

References

- [1] Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. Scape: Shape completion and animation of people. *ACM Trans. Graph.*, 24(3):408–416, jul 2005. 4
- [2] Benjamin Biggs, Ollie Boyne, James Charles, Andrew Fitzgibbon, and Roberto Cipolla. Who left the dogs out: 3D animal reconstruction with expectation maximization in the loop. In *ECCV*, 2020. 2, 3, 4, 6
- [3] Benjamin Biggs, David Novotny, Sebastien Ehrhardt, Hanbyul Joo, Ben Graham, and Andrea Vedaldi. 3d multi-bodies: Fitting sets of plausible 3d human models to ambiguous image data. In *Advances in neural information processing systems*, 2020. 3
- [4] Benjamin Biggs, Thomas Roddick, Andrew Fitzgibbon, and Roberto Cipolla. Creatures great and small: Recovering the shape and motion of animals from video. In *Computer Vision—ACCV 2018: 14th Asian Conference on Computer Vision, Perth, Australia, December 2–6, 2018, Revised Selected Papers, Part V 14*, pages 3–19. Springer, 2019. 3
- [5] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J Black. Keep it simple: Automatic estimation of 3d human pose and shape from a single image. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part V 14*, pages 561–578. Springer, 2016. 3, 4
- [6] Jinkun Cao, Hongyang Tang, Hao-Shu Fang, Xiaoyong Shen, Cewu Lu, and Yu-Wing Tai. Cross-domain adaptation for animal pose estimation. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019. 2, 3, 4
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 4, 7
- [8] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, pages 1180–1189. PMLR, 2015. 3
- [9] Stuart Geman. Statistical methods for tomographic image reconstruction. *Bulletin of International Statistical Institute*, 4:5–21, 1987. 5
- [10] Riza Alp Guler and Iasonas Kokkinos. HoloPose: Holistic 3D human reconstruction in-the-wild. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10884–10894, 2019. 3
- [11] Ju He, Shuo Yang, Shaokang Yang, Adam Kortylewski, Xiaoding Yuan, Jie-Neng Chen, Shuai Liu, Cheng Yang, Qihang Yu, and Alan Yuille. Partimagenet: A large, high-quality dataset of parts. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part VIII*, pages 128–145. Springer, 2022. 1, 4, 6
- [12] Yinghao Huang, Federica Bogo, Christoph Lassner, Angjoo Kanazawa, Peter V. Gehler, Javier Romero, Ijaz Akhter, and Michael J. Black. Towards accurate marker-less human shape and pose estimation over time. In *2017 international conference on 3D vision (3DV)*, 2017. 4
- [13] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE transactions on pattern analysis and machine intelligence*, 36(7):1325–1339, 2013. 4, 7
- [14] Hanbyul Joo, Tomas Simon, and Yaser Sheikh. Total capture: A 3D deformation model for tracking faces, hands, and bodies. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8320–8329, 2018. 4
- [15] Daniel Joska, Liam Clark, Naoya Muramatsu, Ricardo Jericevich, Fred Nicolls, Alexander Mathis, Mackenzie W. Mathis, and Amir Patel. Acinuset: A 3d pose estimation dataset and baseline models for cheetahs in the wild, 2021. 2
- [16] Angjoo Kanazawa, Michael J. Black, David W. Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7122–7131, 2018. 1, 3, 6
- [17] Muhammed Kocabas, Chun-Hao P Huang, Otmar Hilliges, and Michael J Black. Pare: Part attention regressor for 3d human body estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11127–11137, 2021. 1, 6
- [18] Nikos Kolotouros, Georgios Pavlakos, Michael J Black, and Kostas Daniilidis. Learning to reconstruct 3d human pose and shape via model-fitting in the loop. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2252–2261, 2019. 3
- [19] Nikos Kolotouros, Georgios Pavlakos, Michael J. Black, and Kostas Daniilidis. Learning to reconstruct 3D human pose and shape via model-fitting in the loop. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2252–2261, 2019. 3
- [20] Nikos Kolotouros, Georgios Pavlakos, Dinesh Jayaraman, and Kostas Daniilidis. Probabilistic modeling for human mesh recovery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11605–11614, 2021. 3
- [21] Nilesh Kulkarni, Abhinav Gupta, David F Fouhey, and Shubham Tulsiani. Articulation-aware canonical surface mapping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 452–461, 2020. 3
- [22] Christoph Lassner, Javier Romero, Martin Kiefel, Federica Bogo, Michael J. Black, and Peter V. Gehler. Unite the People: Closing the loop between 3D and 2D human representations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4704–4713, 2017. 4
- [23] Chen Li and Gim Hee Lee. From synthetic to real: Unsupervised domain adaptation for animal pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1482–1491, June 2021. 3

- [24] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014. 1
- [25] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014. 4, 7
- [26] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. In *ACM Trans. Graphics (Proc. SIGGRAPH Asia)*, 2015. 1, 4, 7
- [27] Alexander Mathis, Thomas Biasi, Steffen Schneider, Mert Yuksekgonul, Byron Rogers, Matthias Bethge, and Mackenzie W Mathis. Pretraining boosts out-of-domain robustness for pose estimation. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 1859–1868, 2021. 2
- [28] Dushyant Mehta, Helge Rhodin, Dan Casas, Pascal Fua, Oleksandr Sotnychenko, Weipeng Xu, and Christian Theobalt. Monocular 3d human pose estimation in the wild using improved cnn supervision. In *2017 international conference on 3D vision (3DV)*, pages 506–516. IEEE, 2017. 3, 7
- [29] Dushyant Mehta, Helge Rhodin, Dan Casas, Pascal Fua, Oleksandr Sotnychenko, Weipeng Xu, and Christian Theobalt. Monocular 3d human pose estimation in the wild using improved cnn supervision. In *2017 International Conference on 3D Vision (3DV)*, pages 506–516, 2017. 4
- [30] Jiteng Mu, Weichao Qiu, Gregory D Hager, and Alan L Yuille. Learning from synthetic animals. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12386–12395, 2020. 3
- [31] Jiteng Mu, Weichao Qiu, Gregory D. Hager, and Alan L. Yuille. Learning from synthetic animals. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 6
- [32] Mohamed Omran, Christoph Lassner, Gerard Pons-Moll, Peter V. Gehler, and Bernt Schiele. Neural body fitting: Unifying deep learning and model-based human pose and shape estimation. In *2018 international conference on 3D vision (3DV)*, 2018. 3
- [33] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. Osman, Dimitrios Tzionas, and Michael J. Black. Expressive body capture: 3D hands, face, and body from a single image. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10975–10985, 2019. 3, 4
- [34] Georgios Pavlakos, Xiaowei Zhou, Konstantinos G Derpanis, and Kostas Daniilidis. Coarse-to-fine volumetric prediction for single-image 3D human pose. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2017. 3
- [35] Georgios Pavlakos, Luyang Zhu, Xiaowei Zhou, and Kostas Daniilidis. Learning to estimate 3D human pose and shape from a single color image. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 459–468, 2018. 3
- [36] Gregory Rogez, Philippe Weinzaepfel, and Cordelia Schmid. Lcr-net: Localization-classification-regression for human pose. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3433–3441, 2017. 3
- [37] Nadine Rueegg, Silvia Zuffi, Konrad Schindler, and Michael J. Black. Barc: Learning to regress 3d dog shape from images by exploiting breed information. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 3
- [38] Xiao Sun, Bin Xiao, Fangyin Wei, Shuang Liang, and Yichen Wei. Integral human pose regression. In *Proceedings of the European conference on computer vision (ECCV)*, pages 529–545, 2018. 3
- [39] Jun Kai Vince Tan, Ignas Budvytis, and Roberto Cipolla. Indirect deep structured learning for 3d human body shape and pose prediction. 2017. 3
- [40] Hsiao-Yu Tung, Hsiao-Wei Tung, Ersin Yumer, and Katerina Fragkiadaki. Self-supervised learning of motion capture. In *Advances in neural information processing systems*, pages 5236–5246, 2017. 3
- [41] Timo von Marcard, Roberto Henschel, Michael J. Black, Bodo Rosenhahn, and Gerard Pons-Moll. Recovering accurate 3d human pose in the wild using imus and a moving camera. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018. 4
- [42] Donglai Xiang, Hanbyul Joo, and Yaser Sheikh. Monocular total capture: Posing face, body, and hands in the wild. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10965–10974, 2019. 3, 4
- [43] Yu Xiang, Roozbeh Mottaghi, and Silvio Savarese. Beyond pascal: A benchmark for 3d object detection in the wild. In *IEEE Winter Conference on Applications of Computer Vision*, pages 75–82, 2014. 4
- [44] Yuanlu Xu, Song-Chun Zhu, and Tony Tung. Denserac: Joint 3d pose and shape estimation by dense render-and-compare. In *Proceedings of the IEEE/CVF international conference on computer vision*, October 2019. 3
- [45] Chun-Han Yao, Wei-Chih Hung, Yuanzhen Li, Michael Rubinstein, Ming-Hsuan Yang, and Varun Jampani. Lassie: Learning articulated shapes from sparse image ensemble via 3d part discovery. *arXiv preprint arXiv:2207.03434*, 2022. 1, 3
- [46] Hang Yu, Yufei Xu, Jing Zhang, Wei Zhao, Ziyu Guan, and Dacheng Tao. Ap-10k: A benchmark for animal pose estimation in the wild. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021. 2, 4
- [47] Xingyi Zhou, Qixing Huang, Xiao Sun, Xiangyang Xue, and Yichen Wei. Towards 3d human pose estimation in the wild: a weakly-supervised approach. In *Proceedings of*

the IEEE/CVF international conference on computer vision, pages 398–407, 2017. [3](#)

- [48] Silvia Zuffi, Angjoo Kanazawa, Tanya Berger-Wolf, and Michael J. Black. Three-d safari: Learning to estimate zebra pose, shape, and texture from images ”in the wild”. In *International Conference on Computer Vision*, Oct. 2019. [3](#)
- [49] Silvia Zuffi, Angjoo Kanazawa, and Michael J. Black. Lions and tigers and bears: Capturing non-rigid, 3d, articulated shape from images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. [3](#), [6](#)
- [50] Silvia Zuffi, Angjoo Kanazawa, David W Jacobs, and Michael J Black. 3d menagerie: Modeling the 3d shape and pose of animals. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6365–6373, 2017. [1](#), [2](#), [5](#)
- [51] Silvia Zuffi, Angjoo Kanazawa, David W. Jacobs, and Michael J. Black. 3d menagerie: Modeling the 3d shape and pose of animals. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. [3](#), [4](#)