

Prior-Aware Synthetic Data to the Rescue: Animal Pose Estimation with Very Limited Real Data

Le Jiang

Shuangjun Liu

Xiangyu Bai

Sarah Ostadabbas

<https://web.northeastern.edu/ostadabbas>

Augmented Cognition Lab,
Electrical and Computer Engineering
Department, Northeastern University,
Boston, MA, USA

Abstract

Accurately annotated image datasets are essential components for studying animal behaviors from their poses. Compared to the number of species we know and may exist, the existing labeled pose datasets cover only a small portion of them, while building comprehensive large-scale datasets is prohibitively expensive. Here, we present a very data efficient strategy targeted for pose estimation in quadrupeds that requires only a small amount of real images from the target animal. It is confirmed that fine-tuning a backbone network with pretrained weights on generic image datasets such as ImageNet can mitigate the high demand for target animal pose data and shorten the training time by learning the prior knowledge of object segmentation and keypoint estimation in advance. However, when faced with serious data scarcity (i.e., $< 10^2$ real images), the model performance stays unsatisfactory, particularly for limbs with considerable flexibility and several comparable parts. We therefore introduce a prior-aware synthetic animal data generation pipeline called PASyn to augment the animal pose data essential for robust pose estimation. PASyn generates a probabilistically-valid synthetic pose dataset, SynAP, through training a variational generative model on several animated 3D animal models. In addition, a style transfer strategy is utilized to blend the synthetic animal image into the real backgrounds. We evaluate the improvement made by our approach with three popular backbone networks and test their pose estimation accuracy on publicly available animal pose images as well as collected from real animals in a zoo¹.

1 Introduction

The research on animal pose estimation has grown in recent years, covering 2D/3D pose estimation, animal model recovery, and multi-animal pose estimation [13, 15, 16, 17, 19]. The

© 2022. The copyright of this document resides with its authors.

It may be distributed unchanged freely in print or electronic forms.

¹The SynAP dataset and PASyn code is available at <https://github.com/ostadabbas/Prior-aware-Synthetic-Data-Generation-PASyn->.




Figure 1: The outcome samples show the effect of model trained with or without synthetic animal pose (SynAP) dataset. The left side is the pose estimation results based on the DeepLabCut [15] pre-trained on ImageNet and fine-tuned by 99 real animal images and the right side shows the results of the same model when trained on the same amount of real images in addition to our SynAP dataset. Wrong predictions are marked by red cross.

makeup of the available training sets for the animal pose estimation models divides the study into two branches. The first proposes training the model with significant amounts of labeled real data of a single species. The other one is based on a small amount of real data of a target animal and more data from other adjacent domains to make up for the data scarcity of the target animal. Current data scarcity solution is centered around learning animals' common prior knowledge from large amounts of real data of similar species or even humans. Previous research has shown that quadrupeds, including humans, have similar appearances and skeletons, and that knowledge can be shared [2, 14, 20]. This method is easier to implement after the largest labeled dataset for general animal pose estimation, AP10K [27], was made publicly available in 2021. However, there are restrictions imposed by pose label divergence and the high cost of customization in the existing datasets. In contrast to human, there are huge varieties between animals with different length of bones, number of joints, and extra body parts. Large animal datasets with uniform labeling may leave out the special needs of researchers looking for creatures with unique physical traits.

Synthetic data is another potential choice of adjacent domain for the target animal when aiming for both "inexpensive" and "personalization" aspects. Once the synthetic animal model is built, the label can be defined on the target model and annotating pose data becomes considerably faster and less expensive. In addition to synthetic data, previous works [14, 27] have proven that the backbone networks pre-trained on large general datasets such as ImageNet [6] will gain the prior of object segmentation and keypoint prediction. Thus, training even with few frames from the target leads to a high precision pose prediction. However, when there are far fewer images of the target available (e.g. less than 100), the size of the training set is insufficient for a robust model fine-tuning. The accuracy of the model would decrease significantly when facing more actions from free-ranging behaviors, more self or environmental occlusions, and changes in the environment, textures, and shapes. Fig. 1 shows that how EfficientNet-B6 [15] pre-trained on ImageNet and fine-tuned by 99 real animal images is error prone when tested on the images in the wild.

To account for severe data scarcity and guarantee a high degree of label personalization, we present a cost-effective and generic prior-aware synthetic data generation pipeline, called PASyn, for animals pose estimation tasks. In short, this paper contributions are: (1) designing a novel variational autoencoder (VAE)-based synthetic animal data generation pipeline

to generate probabilistically-valid pose data and verifying its performance on several animal pose test sets; (2) blending the synthetic animal images into real backgrounds through style transfer to mitigate the inconsistency between synthetic and real domains; and (3) building a synthetic animal pose (SynAP) dataset through PASyn, containing 3,000 zebra images and extending it with the 3,600 images of six common quadrupeds including, cows, dogs, sheep, deer, giraffes, and horses to make SynAP+ dataset; and (4) releasing a new pose-labeled dataset of mountain zebras in zoo.

2 Related Works

It has been almost ten years since the early animal pose estimation work introduced in [23]. Yet, the performance of these models is still far inferior to the human pose estimation in terms of accuracy, cross-domain adaptation, and model robustness. This is mainly due to a lack of real-world data, which is the challenge for almost all animal pose estimation work.

2.1 Animal Pose Estimation with Label Scarcity

The most common statement in animal pose estimation articles is the lack of labeled dataset for training, which has been mentioned in many works [2, 9, 13, 14, 16, 26]. Numerable variety of species and subspecies, and considerable differences in physical characteristics and behavior patterns between them cause it difficult to form a labeled dataset with adequate samples. The cross-domain adaptation challenge exacerbates the situation. For each new animal, it is necessary to collect data from scratch and label them, since it is insufficient to learn the prior knowledge exclusively from the data of other animals. [2] proposed a cross-domain adaptive method and a large multi-species dataset, Animal-Pose [13] to learn the shared feature space between human and animals, in order to transfer the prior knowledge between them. [27] built their own large multi-species dataset, called AP10k [27] in order to train a robust and cross-domain model. However, both models are still unable to achieve the same level of pose prediction as animals in-domain when facing unseen species.

2.2 Animal Pose Estimation with Synthetic Data

Synthetic data is a promising substitute for real data in previous works. Coarse prior can be learnt and then it would be used to build pseudo-labels for enormous amounts of unlabeled real animal data [13, 16]. The fly in the ointment is that works such as [13, 16] still use significant amounts of real data (such as TigDog dataset [19]) in training, which may not be possible to access for the unseen species. The work in [28], which focuses on animal model recovery, also gives an extraordinary hint on this problem. It purposed a general 3D animal model (called SMAL) by learning the animal toys' scan, and use the SMAL model and a few real pictures to generate a large amount of realistic synthetic data by adjusting the model's texture, shape, pose, background, camera, and other parameters. They also trained an end-to-end network [29] using only synthetic Grevy's zebra data, which came from direct reconstruction from animal 2D images. However, their results are much worse than the current state-of-the-art in terms of 2D animal pose estimation.

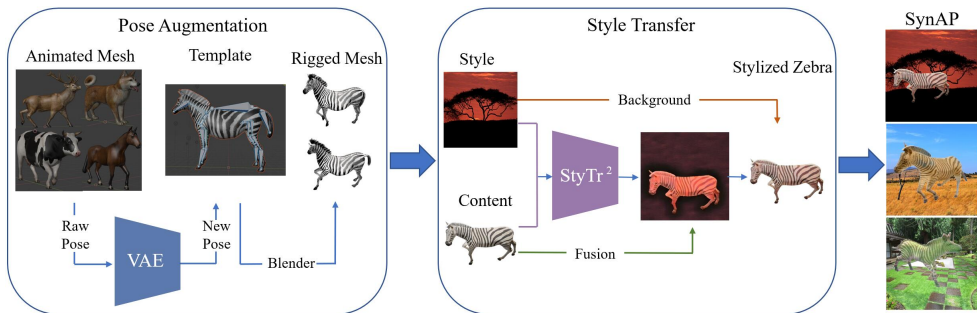


Figure 2: An overview architecture of our prior-aware synthetic data generation (PASyn) pipeline, composed of three parts: pose augmentation, domain transfer and dataset generation. The PASyn pipeline leads to generation of our probabilistically-valid synthetic animal pose (SynAP) dataset.

2.3 Domain Gap between Real and Synthetic

Although synthetic data has advantages in terms of cost, it still faces the problem of fatal domain gap with the real data, which is even challenging to alleviate by current domain adaptation methods [8]. Pasting synthetic animal directly to the real background will result in the strong incongruity, mainly from the difference in projection, brightness, contrast, saturation, etc., between the two images. This incongruity can lead to excessive domain differences between synthetic and real data. Besides, to increase the texture diversity of synthetic animals to enhance the appearance robustness of the model, a common approach is to assign the synthetic data random textures from ImageNet. However, this further increases the discrepancy between synthetic and real domains and thus weakens the improvement of synthetic data in joint training of pose estimation tasks. The approach to increase model robustness and reduce domain variance by style transfer has been applied in medicine [25], monocular depth estimation [10], person re-identification [9]. We, therefore, adopt style transfer to alleviate the domain gap between the synthetic animal and the real background while increasing the texture diversity of the synthetic animal.

3 PASyn Pipeline: Prior-Aware Synthetic Data Generation

Our proposed prior-aware synthetic data generation (PASyn) pipeline enables robust animal pose estimation under severe data scarcity and label divergence. The architecture of PASyn is shown in Fig. 2, where we employ a graphic engine (i.e., Blender) to render the 3D animal mesh into synthetic images. PASyn pipeline is composed of 3 key components: (1) pose data augmentation, (2) synthetic data stylization, and (3) synthetic animal pose image dataset generation. The details are presented in the following subsections. This paper presents PASyn with a working example of zebra as its target animal for pose estimation.

3.1 Capturing Animal Pose Prior

The sample results in Fig. 1 show that when the real data of the target animal is insufficient, the prior knowledge of their pose can be learned from an ingeniously designed synthetic




Figure 3: General quadruped armature. (a) shows the real skeletons of horse (bottom) and dog (top) respectively. (b) shows the artificially designed animal skeleton based on the real one. The red box marks the skeleton of limbs which we are interested in. (c) shows the skeleton of limbs. In (d), v_1 and v_2 represent the spatial orientation of two adjacent bones. The angle between the vectors is the angle between the two bones.

dataset. To learn the animal pose priors, we refer to the variational autoencoder (VAE) framework. VAE has already been proven to be conducive to human pose priors learning [18], when it is trained on large-scale 3D human pose datasets. However, appropriate 3D datasets cannot be easily obtained for animal pose studies. Therefore, we feed several animated low-poly 3D animal models (inexpensively purchased from the CGTrader [9]) to the VAE to learn the probabilistic distribution of the feasible poses. Then, the trained VAE serves as a generative model to create thousands of new poses which are used to rig the template animal meshes in Blender. As seen in Fig. 3(a), the quadrupeds like horses and dogs are similar in their skeleton. Due to the similarity across quadrupeds, the structure Fig. 3(b) is applicable to learning general animal pose priors. Also, considering legs’ considerable flexibility, we only focus on learning the priors of the legs in this work to simplify the model. In order to minimize the influence of skeleton sizes or bone lengths on training, angles between neighboring bones (12 of them) shown in Fig. 3(c) are chosen as primary data for training animal leg pose priors.

Our network follows the basic VAE process, including input, encoder, random sampling, decoder, and output as demonstrated in Fig. 4. $A \in \mathbb{R}^{36}$ is a 1×36 vector containing the angles between 12 pairs of adjacent bones, and each space angle is decomposed into three directions, and \hat{A} is a vector of the same shape representing the generated angles as the model output. The role of the encoder which is composed of dense layers with rectified linear unit (ReLU) activations is to calculate the mean μ and variance σ of the network input. We use the reparameterization trick to randomly sample z in latent space $Z \in \mathbb{R}^{16}$, where $z = \mu + \varepsilon \times \sigma$ and ε satisfies the distribution in the shape of $\mathcal{N}(0, I)$. Finally, the decoder which has a mirror structure of encoder is used to reconstruct the set of space angle, and make the reconstruction results \hat{A} as close to A as possible. Considering that one infeasible angle can ruin an entire pose, during the pre-test, we first assign random sampling from a normal distribution to the variable z , and then we create the histograms of the 12 joint angles with enough \hat{A} ($\approx 10,000$ poses). Based on the statistical results, we can establish the sampling range for each angle. Each range is designed to remove the small amount of angles, which are far from the mean value, and all of the ranges are recorded. After the pre-test, when we generate the new poses, a ‘pose filter’ serves to remove the pose \hat{A} as long as one of the 12 angles is outside these specified ranges and finally produce the refined poses \hat{A}' . We can also obtain poses with higher diversity by increasing the variance of the sampling distribution




Figure 4: An overview architecture of our VAE-based animal pose generative model, composed of two main parts: training (blue-dot box) and test (green-dot box). Before we run “test” to generate new poses, pre-test (brown-dot box) should be performed independently to specify the sampling value ranges for each angles in the ‘pose filter’.

after determining the pose filter. The training loss of the VAE is:

$$\begin{aligned} \mathcal{L}_{\text{total}} &= w_1 \mathcal{L}_{KL} + w_2 \mathcal{L}_{\text{rec}}, \\ \mathcal{L}_{KL} &= KL(q(z|A) || \mathcal{N}(0, I)), \text{ and } \mathcal{L}_{\text{rec}} = \|A - \hat{A}\|_2^2, \end{aligned} \quad (1)$$

where w_i is the weight of each loss term. The Kullback-Leibler term, \mathcal{L}_{KL} , represent the divergence between the encoder’s distribution $q(z|A)$ and $\mathcal{N}(0, I)$. \mathcal{L}_{rec} is the reconstruction term. \mathcal{L}_{KL} encourages a normal distribution noise while \mathcal{L}_{rec} , in contrast, encourages to reconstruct the A without any divergence.

3.2 Stylization: Blending into the Background

As shown in Fig. 2, we use the large number of new synthetic poses obtained through VAE to rig the 3D zebra meshes and enrich synthetic data through changing the ambient light and camera angles. Then, we assign a real (and context-related) background to each rendered zebra. To alleviate the domain gap between the background image I_{sty} (style) and synthetic animal image I_{cont} (content), while increasing the texture diversity of the synthetic animal we employ a style transfer technique. Unlike common convolution-only style transfer methods, we adopt an innovative image style transfer method, called StyTr² [17] based on multi-layer transformers. The I_{sty} and I_{cont} images are segmented into patches to generate feature embedding and eventually formed encoded content sequence by adding positional embedding for each patch. This process is calculated through the content-aware positional encoding based on the semantics of the image content, thereby eliminating the negative effect of scaling on the spatial relationship between patches. Two transformers are applied to encode the content sequence and style sequence, respectively. Then, a transformer decoder is utilized to decode the encoded content sequence according to the encoded style sequence in a regressive fashion. Finally a convolution-based decoder is used to decode the output sequence, consisting of a linear combination of encoded content and style sequence as well as positional embedding, to obtain the stylized content image $I_{\text{cont}}^{\text{sty}}$. Due to strong representation ability of transformers, this method can better capture accurate content representations and avoid loss of details than classical methods.

However, StyTr² cannot freely adjust the intensity of style transfer, unlike the adaptive instance normalization (AdaIN) [18]. Besides StyTr², we also perform a simple pixel summation with normalization of the stylized synthetic animal data with the original data to

control the intensity of style transfer with the fusion rate α . The final stylized synthetic animal image is formulated as $I_{cont}^{final} = (1 - \alpha)I_{cont} + \alpha I_{cont}^{sty}$.

3.3 Creating Synthetic Animal Pose (SynAP) Dataset

Synthetic Animal Pose (SynAP) dataset contains 3,000 synthetic zebra images generated by our prior-aware synthetic animal (PASyn) pipeline. Since zebra was selected as the main species for quantitative evaluation in this work, we collected 5 zebra toys and textured zebra 3D models and synthesized the images from these models with Blender. 3D synthetic model generation pipeline (3D-SMG) introduced in [24] was used to reconstruct two models from zebra toys, while the remaining 3 are adapted from 3D models. In addition to SynAP, which only contains zebra data, SynAP+ extends the SynAP with 3,600 images of 6 other animals, including horses, cows, sheep, dogs, giraffes and deer. We purchased one textured 3D model for each animal from a 3D model website, CGTrader and rendered 600 images for each of the aforementioned animals. The VAE model generates the animal pose in each image and 300 grass, savanna, and forest real scenes are collected from Internet to stylize the synthetic animal. The annotations will be automatically created with each image through calculating the coordinates of each zebra joint in the pose using bone vector transformations.

4 Experimental Analysis

We first describe implementation details of PASyn pipeline and different backbone models which we use to prove its effectiveness in animal pose estimation. In order to achieve a high prediction accuracy on an “unseen” animal in the existing labeled dataset with only a small amount of real data available, we select **zebra** as the main animal to verify the general effect of our PASyn pipeline on different models and quantify the result with the metric, PCK@0.05. we also perform an ablation study on each part of PASyn pipeline and finally show our model’s generalization capacity on different species including deer, cows, dogs, sheep, horses, and giraffes.

4.1 Implementation Details

PASyn Pipeline: We keep the setting in VPoser [18] for training the VAE model. The learning rate of Adam optimizer is 0.001 and we set w_1 and w_2 which are the weights of \mathcal{L}_{KL} and \mathcal{L}_{rec} in Equation (1), as 0.005 and 0.01, respectively. The model is trained on 1,000 poses of animated realistic 3D models for 250 epochs with 128 poses in a batch. The training poses are extracted from the animated horse, dog, sheep, cow, and deer 3D models. Each of these animals has over 20 common actions designed by 3D animators, such as running, walking, and jumping. To increase the diversity of the poses, we choose to generate random samples from a Gaussian distribution $\mathcal{N}(0, 2I)$. The sampling value range of the pose filter will be described in the 4.2. For our domain transfer part, we use the pre-trained transformer decoder provided by [7] and set the fusion rate as 0.5.

Backbone Pose Models: We employ several state-of-the-art pose estimation structures with varying complexity as our backbone networks, including the ResNet-50 and HRNet-W32 of AP10K [27] and EfficientNet-B6 of DeepLabCut [24] to reflect the general effect of our PASyn pipeline. All of the backbones are pre-trained on ImageNet [6].




Figure 5: The histograms of angle values of 12 joints, which are shoulder_right, elbow_right, front-paw_right, shoulder_left, elbow_left, front-paw_left, hip_right, knee_right, back-paw_right, hip_left, knee_left, back-paw_left.

4.2 Pose Filter

Based on the Fig. 5, which is made by decoding of 10,000 random samples from a normal Gaussian distribution, we can set a special sampling range for each joint. The angle value within this range can be regarded as a valid angle, and the pose that satisfies the twelve ranges can be seen as an appropriate pose. These ranges are shoulder_right [40, 100], elbow_right [-125, 0], front-paw_right [-25, 100], shoulder_left [40, 100], elbow_left [-125, 0], front-paw_left [-25, 100], hip_right [-120, -60], knee_right [0, 80], back-paw_right [-125, 0], hip_left [-120, -60], knee_left [0, 80], back-paw_left [-125, 0]. In order to increase the variety of poses and generate more poses with angles near the boundary, we choose to generate random samples from a Gaussian distribution $\mathcal{N}(0, 2I)$ after setting the pose filter. The dropout rate of the pose is 68.0%.

4.3 Evaluation Datasets

Zebra-300 Dataset is the primary test set we use. It contains 40 images from the AP10K test set, 160 unlabeled images randomly selected from AP10K, and 100 images from the Grevy’s zebra dataset [24]. We labeled the images of the latter two according to the AP10K dataset. The images in Zebra-300 dataset are mostly taken in the wild, so the environment occlusion makes the pose estimation task challenging.

Zoo Zebra Dataset is made up of pictures and videos we took at a zoo. The dataset contains 100 zebra images, including 2 different mountain zebra individuals. We performed manual animal pose labeling on the cropped and resized images following the label setting of the AP10K dataset. Each label has 17 key points: nose, eyes, neck, shoulders, elbows, front paws, hips, knees, back paws and root of tail. The occlusion of barbed wire makes this dataset more challenging. We will release this dataset along with the paper.

Table 1: The effect of SynAP with limited real data of pose estimation results of three common backbones, HRNet-w32 [17], EfficientNet-B6 [22], and ResNet-50 [16] tested on Zebra-300 (300 real images). The training set contains only 99 real zebra images (from AP10K) and augmented with SynAP (3000 synthetic zebra images) or SynAP+ (3000 zebra and 2000 other animal synthetic images). Best results for each backbone are shown in bold.

Method	Backbone	Training Set	PCK@0.05 Pose Estimation Accuracy on Zebra-300 Set										
			Eye	Nose	Neck	Shoulders	Elbows	F-Paws	Hips	Knees	B-Paws	RoT	Average
MMPose [15]	HRNet-w32	R(99)	97.3	95.8	83.2	78.8	77.1	62.6	86.0	74.9	59.8	82.4	78.7
		R(99)+ S(3K)	97.8	98.3	81.1	94.0	93.5	92.0	93.7	93.5	89.0	87.6	92.4
		R(99)+ S(5K)	97.5	96.9	81.8	89.6	91.3	90.7	94.1	94.1	90.4	86.0	91.6
DeepLabCut [14]	EfficientNet-B6	R(99)	93.7	96.2	82.5	91.4	80.8	67.4	88.1	84.5	71.8	83.2	83.6
		R(99)+ S(3K)	95.1	97.9	81.5	90.1	83.3	75.5	93.2	89.3	83.9	86.8	87.6
		R(99)+ S(5K)	94.1	92.6	80.8	90.8	87.0	85.7	90.5	93.3	88.3	86.8	89.2
MMPose [15]	ResNet-50	R(99)	96.2	96.9	80.8	59.0	71.3	71.2	88.5	78.2	59.3	85.2	76.9
		R(99)+ S(3K)	95.6	95.8	69.9	87.3	84.6	84.3	90.8	91.2	84.4	77.2	86.7
		R(99)+ S(5K)	97.0	97.9	74.5	85.3	84.2	84.5	94.6	91.4	88.0	85.6	88.4

4.4 Pose Estimation Results

The number of real (R) and synthetic (S) data shown in training set in Table 1 means the total number of images we used for model training. SynAP and SynAP+ are divided into training and validation sets with a ratio of 7 to 1 as the AP10K setting. For real data, considering the data scarcity, we divide it with a ratio of 4 to 1. From Table 1, we can clearly know that only using 99 images for training is insufficient for any backbone network. The highest prediction accuracy is 83.2%, reached by EfficientNet-B6. After adding SynAP to the training set, the prediction results dramatically increase to around 90%. And after adding SynAP+ to the training set, the accuracy is further improved. Among them, HRNet-w32 achieves the highest accuracy of 92.4% when trained with SynAP.

Different animals show domain similarity and the possibility of mutual transfer learning, which is proven in many works [9, 22]. Thus, we want to show that even if we don't ignore the existing labeled data, our SynAP dataset can still improve the prediction of the unseen animals. Table 1 and Table 2 show that training the model with 99 zebra images and 8,000 images of the other animals can increase the average PCK from 78.7% to 91.4%. However, Table 1 shows that even with only 99 real we are still able to surpass that result and reach SOTA by adding SynAP and SynAP+. Moreover, Table 2 points that the performance of the models trained with 8,000 real data in addition to SynAP or SynAP+ is further improved to 93.8% and 94.2%. When there is no target animal (i.e., zebra) in the training set, while the state-of-the-art model suffers from this situation (accuracy drops to 78.3%), the model can recover its performance if given the SynAP dataset, as Table 2 shows. This confirms that our method can achieve high-precision prediction even for the unseen animals.

4.5 Ablation Study

As seen from Table 1, even if it is only trained on a small amount of real data, the model can accurately predict the keypoints with apparent features and less flexibility, such as nose and eyes. Various alternative candidate positions and large degrees of freedom of keypoints on the limbs would considerably affect the prediction results. The main improvement brought by SynAP and SynAP+ is the significant increase of the prediction of keypoints in animal limbs. To verify that the VAE, StyTr² models [12], and the higher variance of VAE sampling distribution in our PASyn pipeline can mitigate the mismatching prediction of the limbs and reduce the domain discrepancy between the real (R) and synthetic (S) domains, we test them on the Zoo Zebra and Zebra-300 datasets. Table 3 shows the outcomes of eight experiments (a-h) with various training conditions. When the model trained without VAE, we set the

Table 2: The effect of SynAP with large real data on HRNet-w32 [21] tested on Zebra-300 (300 real images). The training set contains 8000 images from real animals (from AP10K) and augmented with SynAP (3000 synthetic zebra images) or SynAP+ (3000 zebra and 2000 other animal synthetic images). We tested the results when the 8000 images contain 99 real zebra inside them and when no zebra images is used. Please note that the HRNet-w32 model trained on AP10K is the SOTA. Best results are shown in bold.

Training Set	Real Zebra	PCK@0.05 Pose Estimation Accuracy on Zebra-300 Set										
		Eye	Nose	Neck	Shoulders	Elbows	F-Paws	Hips	Knees	B-Paws	RoT	Average
R(8K) (SOTA)	✓	97.5	97.2	79.4	87.8	90.3	93.8	95.3	94.1	89.5	86.4	91.4
R(8K) + S(3K)		97.3	98.3	79.0	93.1	94.9	96.0	95.3	96.7	93.3	89.6	93.8
R(8K) + S(5K)		97.3	97.6	81.1	93.7	95.7	96.0	96.6	96.0	94.3	87.6	94.2
R(8K)	✗	79.7	87.7	37.4	77.6	80.0	87.6	82.0	86.4	81.3	67.2	78.3
R(8K) + S(3K)		94.8	96.2	67.1	90.8	87.9	90.2	87.6	91.6	89.7	77.6	88.2
R(8K) + S(5K)		97.5	96.2	66.1	91.6	89.5	93.8	93.9	93.5	91.1	77.6	90.2

Table 3: Ablation study on two parts of PASyn pipeline, VAE, style transfer and σ^2 random sampling distribution on two different zebra datasets with resolution of 300×300 . The HRNet-w32 is the backbone here. The training set contains 99 real zebra (from AP10K) and augmented with SynAP (3000 synthetic zebra images). Best results are shown in bold.

Index	Training Set	VAE	Style Transfer	σ^2	Zoo Zebra Set	Zebra-300 Set
a	R(99)	✗	✗	✗	76.0	78.7
b	S(3K)	✗	✗	21	38.7	30.0
c	S(3K)	✓	✗	21	44.2	36.7
d	S(3K)	✓	✓	21	42.9	46.6
e	R(99)+S(3K)	✗	✗	21	89.8	88.0
f	R(99)+S(3K)	✓	✗	21	90.4	89.8
g	R(99)+S(3K)	✓	✓	1	90.5	91.1
h	R(99)+S(3K)	✓	✓	21	91.5	92.4

range according to the animal pose/animation data we have and uniformly sample the pose data from this range, and conduct randomly sampling, similar to the method used in [16, 29]. We will keep the original texture of the model when StyTr² is not used. The comparisons of (b), (c), (d), and (e), (f), (h) clearly reflect that the model trained with both of VAE and StyTr² in our PASyn pipeline can achieve obviously higher accuracy than the model trained without them. Moreover, the comparison of (g) and (h) shows how our designed value range affects the prediction results. In (g), although VAE is used, we only choose to generate random samples from a normal distribution $\mathcal{N}(0, I)$, and do not use the pose filter. On the other hand, in (h), we increase the variance of the Gaussian distribution from 1 to 2 to enrich the generated pose, and at the same time, we use the pose filter to filter out the poses with extreme angles caused by increasing the variance. The experimental results also prove that our proposed method is better than the traditional VAE strategy.

5 Conclusion

We presented a cost-effective and general prior-aware synthetic data generation pipeline (PASyn) for animals pose studies that suffer from a severe data scarcity. A probabilistically-valid variational generative model as well as a style transfer strategy are introduced to increase the validity of the generated poses and to reduce the domain discrepancy between the synthetic and real images. The positive effect of our synthetic animal pose dataset SynAP and its extended version SynAP+ on pose estimation task of animals with a small amount of real data is verified on different backbones and achieves state-of-the-art performance.

References

- [1] Amir Atapour-Abarghouei and Toby P. Breckon. Real-time monocular depth estimation using synthetic data with domain adaptation via image style transfer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [2] Jinkun Cao, Hongyang Tang, Hao-Shu Fang, Xiaoyong Shen, Cewu Lu, and Yu-Wing Tai. Cross-domain adaptation for animal pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9498–9507, 2019.
- [3] CGTrader. The world’s preferred source for 3d content. <https://www.cgtrader.com/>, 2022.
- [4] Yanwen Chong, Chengwei Peng, Jingjing Zhang, and Shaoming Pan. Style transfer for unsupervised domain-adaptive person re-identification. *Neurocomputing*, 422:314–321, 2021.
- [5] MMPose Contributors. Openmmlab pose estimation toolbox and benchmark. <https://github.com/open-mmlab/mmpose>, 2022.
- [6] 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.
- [7] Yingying Deng, Fan Tang, Weiming Dong, Chongyang Ma, Xingjia Pan, Lei Wang, and Changsheng Xu. Stytr2: Image style transfer with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11326–11336, 2022.
- [8] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The journal of machine learning research*, 17(1):2096–2030, 2016.
- [9] Jacob M Graving, Daniel Chae, Hemal Naik, Liang Li, Benjamin Koger, Blair R Costelloe, and Iain D Couzin. Deepposekit, a software toolkit for fast and robust animal pose estimation using deep learning. *Elife*, 8:e47994, 2019.
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [11] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE international conference on computer vision*, pages 1501–1510, 2017.
- [12] Cao Jinkun, Tang Hongyang, Fang Hao-Shu, Shen Xiaoyong, Lu Cewu, and Tai Yu-Wing. Animal pose dataset. <https://sites.google.com/view/animal-pose/>, 2019.
- [13] 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*, pages 1482–1491, 2021.

- [14] Alexander Mathis, Pranav Mamidanna, Kevin M Cury, Taiga Abe, Venkatesh N Murthy, Mackenzie Weygandt Mathis, and Matthias Bethge. Deeplabcut: markerless pose estimation of user-defined body parts with deep learning. *Nature neuroscience*, 21(9):1281–1289, 2018.
- [15] Alexander Mathis, Thomas Biasi, Steffen Schneider, Mert Yuksekogul, 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 (WACV)*, pages 1859–1868, January 2021.
- [16] 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.
- [17] Natalia Neverova, David Novotny, Vasil Khalidov, Marc Szafraniec, Patrick Labatut, and Andrea Vedaldi. Continuous surface embeddings, 2020.
- [18] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. A. Osman, Dimitrios Tzionas, and Michael J. Black. Expressive body capture: 3d hands, face, and body from a single image. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [19] L Del Pero, Ricco Susanna, Sukthankar Rahul, and Ferrari Vittorio. Tigdog dataset. <http://calvin-vision.net/datasets/tigdog/>, 2016.
- [20] Artsiom Sanakoyeu, Vasil Khalidov, Maureen S. McCarthy, Andrea Vedaldi, and Natalia Neverova. Transferring dense pose to proximal animal classes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [21] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5693–5703, 2019.
- [22] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pages 6105–6114. PMLR, 2019.
- [23] Sara Vicente and Lourdes Agapito. Balloon shapes: Reconstructing and deforming objects with volume from images. In *2013 International Conference on 3D Vision-3DV 2013*, pages 223–230. IEEE, 2013.
- [24] Kathan Vyas, Le Jiang, Shuangjun Liu, and Sarah Ostadabbas. An efficient 3d synthetic model generation pipeline for human pose data augmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1542–1552, 2021.
- [25] Sophia J Wagner, Nadiieh Khalili, Raghav Sharma, Melanie Boxberg, Carsten Marr, Walter de Back, and Tingying Peng. Structure-preserving multi-domain stain color augmentation using style-transfer with disentangled representations. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 257–266. Springer, 2021.

- [26] Anqi Wu, Estefany Kelly Buchanan, Matthew Whiteway, Michael Schartner, Guido Meijer, Jean-Paul Noel, Erica Rodriguez, Claire Everett, Amy Norovich, Evan Schaffer, et al. Deep graph pose: a semi-supervised deep graphical model for improved animal pose tracking. *Advances in Neural Information Processing Systems*, 33:6040–6052, 2020.
- [27] Hang Yu, Yufei Xu, Jing Zhang, Wei Zhao, Ziyu Guan, and Dacheng Tao. Ap-10k: A benchmark for animal pose estimation in the wild. *arXiv preprint arXiv:2108.12617*, 2021.
- [28] 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.
- [29] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5359–5368, 2019.